Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide

Cisco MDS NX-OS Release 6.2(1)
April 2013

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Text Part Number: OL-29284-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide
© 2012-2013 Cisco Systems, Inc. All rights reserved.

New and Changed Information xiii

Preface xv

Audience xv

Organization xv

Document Conventions xvi

Related Documentation xvii

Release Notes xvii

Regulatory Compliance and Safety Information xvii

Compatibility Information xvii

Hardware Installation xvii

Software Installation and Upgrade xvii

Cisco NX-OS xvii

Command-Line Interface xviii

Intelligent Storage Networking Services Configuration Guides xviii

Troubleshooting and Reference xviii

Obtaining Documentation and Submitting a Service Request xviii

CHAPTER 1 **Interfaces Overview** 1-1

Trunks and PortChannels 1-1

Fibre Channel Port Rate Limiting 1-1

Extended Credits 1-2

N Port Virtualization 1-2

FlexAttach 1-2

CHAPTER 2 **Configuring Interfaces** 2-1

Information About Interfaces 2-1

Interface Description 2-2

Interface Modes 2-2

E Port 2-3

F Port 2-3

FL Port 2-3

NP Ports 2-3

TL Port 2-4

TE Port 2-4
TF Port 2-4
TNP Port 2-5
SD Port 2-5
ST Port 2-5
Fx Port 2-5
B Port 2-5
Auto Mode 2-5
Interface States 2-6
Administrative States 2-6
Operational States 2-6
Reason Codes 2-6
Graceful Shutdown 2-9
Port Administrative Speeds 2-10
Autosensing 2-10
Frame Encapsulation 2-10
Beacon LEDs 2-11
Speed LEDs 2-11
Bit Error Thresholds 2-11
SFP Transmitter Types 2-12
TL Ports 2-13
TL Port ALPA Caches 2-14
Port Guard 2-14
Port Monitor 2-15
Port Monitor Port Guard 2-16
Port Group Monitor 2-16
Local Switching 2-16
Slow Drain Device Detection and Congestion Avoidance 2-17
Management Interfaces 2-17
VSAN Interfaces 2-18
Prerequisites for Interfaces 2-18
Guidelines and Limitations 2-18
Generation 1 Interface Configuration Guidelines 2-18
Private Loop Configuration Guidelines 2-19
VSAN Interface Configuration Guidelines 2-19
Default Settings 2-20
Configuring Interfaces 2-20
Configuring Fibre Channel Interfaces 2-21
Setting the Interface Administrative State 2-22

Configuring Interface Modes 2-22
Configuring System Default Port Mode F 2-23
Configuring ISL between Two Switches 2-24
Configuring 10-Gbps FC Mode 2-24
Configuring Port Administrative Speeds 2-25
Configuring Port Speed Group 2-26
Configuring the Interface Description 2-26
Specifying a Port Owner 2-27
Configuring Beacon Mode 2-27
Disabling Bit Error Threshold 2-28
Configuring Switch Port Attribute Default Values 2-28
Configuring TL Ports 2-29
Manually Inserting Entries into the ALPA Cache 2-29
Clearing the ALPA Cache 2-29
Configuring Port Guard 2-30
Configuring Port Monitor 2-31
Enabling Port Monitor 2-31
Configuring a Port Monitor Policy 2-32
Activating a Port Monitor Policy 2-34
Configuring a Port Monitor Port Guard 2-34
Configuring Port Group Monitor 2-34
Enabling Port Group Monitor 2-35
Configuring a Port Group Monitor Policy 2-35
Reverting to the Default Policy for a Specific Counter 2-36
Turning Off the Monitoring of Specific Counter 2-36
Activating a Port Group Monitor Policy 2-37
Configuring Management Interfaces 2-37
Creating VSAN Interfaces 2-38
Configuring Slow Drain Device Detection and Congestion Avoidance 2-38
Configuring Congestion Frame Timeout Value 2-39
Configuring Stuck Frame Timeout Value 2-39
Configuring No-Credit Timeout Value 2-39
Configuring Credit Loss Recovery Threshold and Action 2-40
Configuring Average Credit Nonavailable Duration Threshold and Action 2-41
Verifying Interfaces Configuration 2-41
Displaying Interface Information 2-42
Displaying TL Port Information 2-50
Displaying the ALPA Cache Contents 2-51
Displaying Port Monitor Status and Policies 2-51
Displaying Port Group Monitor Status and Policies 2-53

Displaying Management Interface Configuration 2-55
Displaying VSAN Interface Information 2-55

CHAPTER 3 **Configuring Fibre Channel Interfaces 3-1**

Information About Fibre Channel Interfaces 3-1
Generations of Modules and Switches 3-1
Port Groups 3-3
Port Rate Modes 3-5
Dedicated Rate Mode 3-7
Shared Rate Mode 3-8
Dedicated Rate Mode Configurations for the 8-Gbps Modules 3-9
Port Speed 3-10
Dynamic Bandwidth Management 3-10
Out-of-Service Interfaces 3-11
Oversubscription Ratio Restrictions 3-11
Bandwidth Fairness 3-17
Upgrade or Downgrade Scenario 3-17
Guidelines and Limitations 3-18
Combining Generation 1, Generation 2, Generation 3, and Generation 4 Modules 3-18
Local Switching Limitations 3-19
Port Index Limitations 3-19
PortChannel Limitations 3-21
Default Settings 3-25
Configuring Fibre Channel Interfaces 3-26
Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode 3-26
Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode 3-27
Task Flow for Configuring 12-Port 4-Gbps Module Interfaces 3-28
Task Flow for Configuring 4-Port 10-Gbps Module Interfaces 3-28
Configuring Port Speed 3-29
Configuring Rate Mode 3-30
Displaying the Rate Mode Configuration for Interfaces 3-31
Configuring Local Switching 3-36
Disabling Restrictions on Oversubscription Ratios 3-37
Enabling Restrictions on Oversubscription Ratios 3-39
Enabling Bandwidth Fairness 3-40
Disabling Bandwidth Fairness 3-41
Taking Interfaces Out of Service 3-41
Releasing Shared Resources in a Port Group 3-42
Disabling ACL Adjacency Sharing for System Image Downgrade 3-43

Verifying Fibre Channel Interfaces Configuration 3-43
Displaying Interface Capabilities 3-44
Displaying SFP Diagnostic Information 3-45
Configuration Examples for Fibre Channel Interfaces 3-45
Configuration Example for 48-Port 8-Gbps Module Interfaces 3-46
Configuration Example for 24-Port 8-Gbps Module Interfaces 3-46
Configuration Example for 4/44-Port 8-Gbps Module Interfaces 3-47
Configuration Example for 48-Port 4-Gbps Module Interfaces 3-48
Configuration Example for 24-Port 4-Gbps Module Interfaces 3-49

CHAPTER 4 **Configuring Interface Buffers** 4-1
Information About Interface Buffers 4-1
Buffer-to-Buffer Credits 4-1
Performance Buffers 4-2
Buffer Pools 4-2
BB_Credit Buffers for Switching Modules 4-5
Configuring Buffer Credits on a Generation 2, Generation 3 or Generation 4 Module 4-5
48-Port 8-Gbps Advanced Fibre Channel Module BB_Credit Buffers 4-6
48-Port 8-Gbps Fibre Channel Module BB_Credit Buffers 4-7
24-Port 8-Gbps Fibre Channel Module BB_Credit Buffers 4-8
4/44-Port 8-Gbps Host-Optimized Fibre Channel Module BB_Credit Buffers 4-9
48-Port 4-Gbps Fibre Channel Module BB_Credit Buffers 4-10
24-Port 4-Gbps Fibre Channel Module BB_Credit Buffers 4-12
18-Port Fibre Channel/4-Port Gigabit Ethernet Multiservice Module BB_Credit Buffers 4-13
12-Port 4-Gbps Switching Module BB_Credit Buffers 4-13
4-Port 10-Gbps Switching Module BB_Credit Buffers 4-15
BB_Credit Buffers for Fabric Switches 4-16
Cisco MDS 9148 Fabric Switch BB_Credit Buffers 4-16
Cisco MDS 9134 Fabric Switch BB_Credit Buffers 4-16
Cisco MDS 9124 Fabric Switch BB_Credit Buffers 4-17
Cisco MDS 9222i Multiservice Modular Switch BB_Credit Buffers 4-17
Extended BB_Credits 4-17
Extended BB_credits on Generation 1 Switching Modules 4-18
Extended BB_credits on Generation 2 and Generation 3 Switching Modules 4-19
Buffer-to-Buffer Credit Recovery 4-20
Buffer-to-Buffer State Change Number 4-20
Receive Data Field Size 4-21
Configuring Interface Buffers 4-21
Configuring Buffer-to-Buffer Credits 4-21

Configuring Performance Buffers 4-22
Configuring Extended BB_credits 4-23
Enabling Buffer-to-Buffer Credit Recovery 4-24
Enabling the Buffer-to-Buffer State Change Number 4-24
Configuring Receive Data Field Size 4-24
Verifying BB_Credit Configuration 4-25

CHAPTER 5 **Configuring Trunking** 5-1
Information About Trunking 5-1
Trunking E Ports 5-1
Trunking F Ports 5-2
Key Concepts 5-3
Trunking Protocols 5-3
Trunk Modes 5-4
Trunk-Allowed VSAN Lists and VF_IDs 5-5
Guidelines and Limitations 5-7
General Guidelines and Limitations 5-7
Upgrade and Downgrade Limitations 5-8
Difference Between TE Ports and TF-TNP Ports 5-8
Trunking Misconfiguration Examples 5-10
Default Settings 5-11
Configuring Trunking 5-11
Enabling the Cisco Trunking and Channeling Protocols 5-11
Enabling the F Port Trunking and Channeling Protocol 5-12
Configuring Trunk Mode 5-12
Configuring an Allowed-Active List of VSANs 5-12
Verifying Trunking Configuration 5-13
Configuration Example for F Port Trunking 5-14

CHAPTER 6 **Configuring PortChannels** 6-1
Information About PortChannels 6-1
PortChannels Overview 6-2
E PortChannels 6-2
F and TF PortChannels 6-3
PortChanneling and Trunking 6-3
Load Balancing 6-4
PortChannel Modes 6-6
PortChannel Deletion 6-7
Interfaces in a PortChannel 6-7

Interface Addition to a PortChannel 6-8
Forcing an Interface Addition 6-9
Interface Deletion from a PortChannel 6-9
PortChannel Protocols 6-9
Channel Group Creation 6-10
Autocreation 6-11
Manually Configured Channel Groups 6-12
Prerequisites for PortChannels 6-12
Guidelines and Limitations 6-13
General Guidelines and Limitations 6-13
Generation 1 PortChannel Limitations 6-14
F and TF PortChannel Limitations 6-14
Valid and Invalid PortChannel Examples 6-15
Default Settings 6-16
Configuring PortChannels 6-17
Configuring PortChannels Using the WizardCreating a PortChannel 6-17
Configuring the PortChannel Mode 6-17
Deleting PortChannels 6-18
Adding an Interface to a PortChannel 6-18
Forcing an Interface Addition 6-19
Deleting an Interface from a PortChannel 6-19
Enabling and Configuring Autocreation 6-20
Converting to Manually Configured Channel Groups 6-20
Verifying PortChannel Configuration 6-20
Configuration Examples for F and TF PortChannels 6-24

CHAPTER 7 **Configuring N Port Virtualization** 7-1
Information About N Port Virtualization 7-1
NPV Overview 7-1
N Port Identifier Virtualization 7-2
N Port Virtualization 7-2
NPV Mode 7-4
NP Ports 7-5
NP Links 7-5
Internal FLOGI Parameters 7-5
Default Port Numbers 7-6
NPV CFS Distribution over IP 7-7
NPV Traffic Management 7-7
Auto 7-7

Traffic Map 7-7
Disruptive 7-8
Multiple VSAN Support 7-8
Guidelines and Limitations 7-8
NPV Guidelines and Requirements 7-8
NPV Traffic Management Guidelines 7-9
DPVM Configuration Guidelines 7-9
NPV and Port Security Configuration Guidelines 7-10
Configuring N Port Virtualization 7-10
Enabling N Port Identifier Virtualization 7-10
Configuring NPV 7-10
Configuring NPV Traffic Management 7-12
Configuring List of External Interfaces per Server Interface 7-12
Enabling the Global Policy for Disruptive Load Balancing 7-13
Verifying NPV Configuration 7-13
Verifying NPV 7-14
Verifying NPV Traffic Management 7-15

CHAPTER 8 **Configuring FlexAttach Virtual pWWN** 8-1
Information About FlexAttach Virtual pWWN 8-1
FlexAttach Virtual pWWN 8-1
Difference Between San Device Virtualization and FlexAttach Port Virtualization 8-2
FlexAttach Virtual pWWN CFS Distribution 8-2
Security Settings for FlexAttach Virtual pWWN 8-3
Guidelines and Limitations 8-3
Configuring FlexAttach Virtual pWWN 8-3
Automatically Assigning FlexAttach Virtual pWWN 8-3
Manually Assigning FlexAttach Virtual pWWN 8-4
Mapping pWWN to Virtual pWWN 8-4
Verifying FlexAttach Virtual pWWN Configuration 8-5
Verifying the End Device 8-6
Monitoring FlexAttach Virtual pWWN 8-6

CHAPTER 9 **Configuring Port Tracking** 9-1
Information About Port Tracking 9-1
Guidelines and Limitations 9-2
Default Settings 9-2
Configuring Port Tracking 9-3

Enabling Port Tracking 9-3
Information About Configuring Linked Ports 9-3
Binding a Tracked Port Operationally 9-4
Information About Tracking Multiple Ports 9-4
Tracking Multiple Ports 9-5
Information About Monitoring Ports in a VSAN 9-5
Monitoring Ports in a VSAN 9-5
Information AboutForceful Shutdown 9-6
Forcefully Shutting Down a Tracked Port 9-6
Displaying Port Tracking Information 9-6
9-8

New and Changed Information

As of Cisco MDS NX-OS Release 4.2(1), software configuration information is available in new feature-specific configuration guides for the following information:

- System management
- Interfaces
- Fabric
- Quality of service
- Security
- IP services
- High availability and redundancy

The information in these new guides previously existed in the *Cisco MDS 9000 Family CLI Configuration Guide* and in the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*. Those configuration guides remain available on Cisco.com and should be used for all software releases prior to Fabric Manager Release 5.0(1a). Each guide addresses the features introduced in or available in a particular release. Select and view the configuration guide that pertains to the software installed in your switch.

For a complete list of document titles, see the list of Related Documentation in the "Preface."

http://www.cisco.com/en/US/products/ps5989/prod_release_notes_list.htm

About This Guide

Table 1 lists the New and Changed features for this guide, starting with MDS NX-OS Release 5.2(1).

Table 1 New and Changed Features

Feature	GUI Change	Description	Changed in Release	Where Documented
Port Monitor Enhancements	Configuring Port Monitor Policy	Added information about the feature Port Monitor Port Guard and three new counters for the port monitor command.	5.2(2a)	Chapter 2, "Configuring Interfaces"

Table 1 New and Changed Features (continued)

Feature	GUI Change	Description	Changed in Release	Where Documented
FlexAttach	Disable FlexAttach	Added information about disabling FlexAttach.	5.0(1a)	Chapter 8, "Configuring FlexAttach Virtual pWWN"
Port Group Monitoring Enhancements	Check Oversubscription > Monitor	Added information about monitoring a selected port group.	5.0(1a)	Chapter 2, "Configuring Interfaces"

Preface

This preface describes the audience, organization, and conventions of the *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*. It also provides information on how to obtain related documentation.

Audience

This guide is for experienced network administrators who are responsible for configuring and maintaining the Cisco MDS 9000 Family of multilayer directors and fabric switches.

Organization

The *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide* is organized as follows:

Chapter	Title	Description
Chapter 1	Interfaces Overview	Provides an overview of all the features in this guide.
Chapter 2	Configuring Interfaces	Explains Generation 1 and Generation 2 module port and operational state concepts in Cisco MDS 9000 Family switches and provides details on configuring ports and interfaces.
Chapter 3	Configuring Fibre Channel Interfaces	Explains configuration concepts for Fibre Channel module ports and interfaces.
Chapter 4	Configuring Interface Buffers	Explains configuration concepts for Interface Buffers.
Chapter 5	Configuring Trunking	Explains TE ports and trunking concepts.
Chapter 6	Configuring PortChannels	Explains PortChannels and load balancing concepts and provides details on configuring PortChannels, adding ports to PortChannels, and deleting ports from PortChannels.

Chapter	Title	Description
Chapter 7	Configuring N Port Virtualization	Provides an overview of N Port Virtualization and includes quidelines and requirements for configuring and verifying NPV.
Chapter 8	Configuring FlexAttach Virtual pWWN	FlexAttach virtual pWWN feature facilitates server and configuration management. In a SAN environment, the server installation or replacement, requires interaction and coordination among the SAN and server administrators.

Document Conventions

Command descriptions use these conventions:

boldface font	Commands and keywords are in boldface.
italic font	Arguments for which you supply values are in italics.
[ ]	Elements in square brackets are optional.
[x | y | z]	Optional alternative keywords are grouped in brackets and separated by vertical bars.

Screen examples use these conventions:

`screen font`	Terminal sessions and information the switch displays are in screen font.
`boldface screen font`	Information you must enter is in boldface screen font.
italic screen font	Arguments for which you supply values are in italic screen font.
`< >`	Nonprinting characters, such as passwords, are in angle brackets.
`[ ]`	Default responses to system prompts are in square brackets.
`!, #`	An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.

This document uses the following conventions:

Note Means reader *take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Caution Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Related Documentation

The documentation set for the Cisco MDS 9000 Family includes the following documents. To find a document online, use the Cisco MDS NX-OS Documentation Locator at:

http://www.cisco.com/en/US/docs/storage/san_switches/mds9000/roadmaps/doclocater.htm

Release Notes

- *Cisco MDS 9000 Family Release Notes for Cisco MDS NX-OS Releases*
- *Cisco MDS 9000 Family Release Notes for MDS SAN-OS Releases*
- *Cisco MDS 9000 Family Release Notes for Cisco MDS 9000 EPLD Images*
- *Cisco DCNM Release Notes*

Regulatory Compliance and Safety Information

- *Regulatory Compliance and Safety Information for the Cisco MDS 9000 Family*

Compatibility Information

- *Cisco Data Center Interoperability Support Matrix*
- *Cisco MDS 9000 NX-OS Hardware and Software Compatibility Information and Feature Lists*
- *Cisco MDS 9000 Family Switch-to-Switch Interoperability Configuration Guide*

Hardware Installation

- *Cisco MDS 9500 Series Hardware Installation Guide*
- *Cisco MDS 9200 Series Hardware Installation Guide*
- *Cisco MDS 9100 Series Hardware Installation Guide*
- *Cisco MDS 9124 and Cisco MDS 9134 Multilayer Fabric Switch Quick Start Guide*

Software Installation and Upgrade

- *Cisco MDS 9000 NX-OS Software Upgrade and Downgrade Guide*

Cisco NX-OS

- *Cisco MDS 9000 Family NX-OS Licensing Guide*
- *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Interfaces Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*

- *Cisco MDS 9000 Family NX-OS Quality of Service Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Security Configuration Guide*
- *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Intelligent Storage Services Configuration Guide*
- *Cisco MDS 9000 Family NX-OS High Availability and Redundancy Configuration Guide*
- *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*
- *Cisco MDS 9000 Family Cookbook for Cisco MDS SAN-OS*

Command-Line Interface

- *Cisco MDS 9000 Family Command Reference*

Intelligent Storage Networking Services Configuration Guides

- *Cisco MDS 9000 Family I/O Acceleration Configuration Guide*
- *Cisco MDS 9000 Family SANTap Deployment Guide*
- *Cisco MDS 9000 Family Data Mobility Manager Configuration Guide*
- *Cisco MDS 9000 Family Storage Media Encryption Configuration Guide*

Troubleshooting and Reference

- *Cisco MDS 9000 Family and Nexus 7000 Series System Messages Reference*
- *Cisco MDS 9000 Family SAN-OS Troubleshooting Guide*
- *Cisco MDS 9000 Family NX-OS MIB Quick Reference*
- *Cisco DCNM for SAN Database Schema Reference*

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html

- Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Interfaces Overview

- Trunks and PortChannels, page 1-1
- Fibre Channel Port Rate Limiting, page 1-1
- Extended Credits, page 1-2
- N Port Virtualization, page 1-2
- FlexAttach, page 1-2

Trunks and PortChannels

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports.

PortChannels aggregate multiple physical ISLs into one logical link with higher bandwidth and port resiliency for both Fibre Channel and FICON traffic. With this feature, up to 16 expansion ports (E-ports) or trunking E-ports (TE-ports) can be bundled into a PortChannel. ISL ports can reside on any switching module, and they do not need a designated master port. If a port or a switching module fails, the PortChannel continues to function properly without requiring fabric reconfiguration.

Cisco NX-OS software uses a protocol to exchange PortChannel configuration information between adjacent switches to simplify PortChannel management, including misconfiguration detection and autocreation of PortChannels among compatible ISLs. In the autoconfigure mode, ISLs with compatible parameters automatically form channel groups; no manual intervention is required.

PortChannels load balance Fibre Channel traffic using a hash of source FC-ID and destination FC-ID, and optionally the exchange ID. Load balancing using PortChannels is performed over both Fibre Channel and FCIP links. Cisco NX-OS software also can be configured to load balance across multiple same-cost FSPF routes.

Fibre Channel Port Rate Limiting

The Fibre Channel port rate-limiting feature for the Cisco MDS 9100 Series controls the amount of bandwidth available to individual Fibre Channel ports within groups of four host-optimized ports. Limiting bandwidth on one or more Fibre Channel ports allows the other ports in the group to receive a greater share of the available bandwidth under high-utilization conditions. Port rate limiting is also beneficial for throttling WAN traffic at the source to help eliminate excessive buffering in Fibre Channel and IP data network devices.

Extended Credits

Full line-rate Fibre Channel ports provide at least 255 buffer credits standard. Adding credits lengthens distances for Fibre Channel SAN extension. Using extended credits, up to 4095 buffer credits from a pool of more than 6000 buffer credits for a module can be allocated to ports as needed to greatly extend the distance for Fibre Channel SANs.

N Port Virtualization

Cisco NX-OS software supports industry-standard N port identifier virtualization (NPIV), which allows multiple N port fabric logins concurrently on a single physical Fibre Channel link. HBAs that support NPIV can help improve SAN security by enabling zoning and port security to be configured independently for each virtual machine (OS partition) on a host. In addition to being useful for server connections, NPIV is beneficial for connectivity between core and edge SAN switches.

N port virtualizer (NPV) is a complementary feature that reduces the number of Fibre Channel domain IDs in core-edge SANs. Cisco MDS 9000 family fabric switches operating in the NPV mode do not join a fabric; they only pass traffic between core switch links and end devices, which eliminates the domain IDs for these switches. NPIV is used by edge switches in the NPV mode to log in to multiple end devices that share a link to the core switch. This feature is available only for Cisco MDS Blade Switch Series, the Cisco MDS 9124 Multilayer Fabric Switch, and the Cisco MDS 9134 Multilayer Fabric Switch.

FlexAttach

Cisco NX-OS supports the FlexAttach feature. One of the main problems in a SAN environment is the time and effort required to install and replace servers. The process involves both SAN and server administrators, and the interaction and coordination between them can make the process time consuming. To alleviate the need for interaction between SAN and server administrators, the SAN configuration should not be changed when a new server is installed or an existing server is replaced. FlexAttach addresses these problems, reducing configuration changes and the time and coordination required by SAN and server administrators when installing and replacing servers. This feature is available only for Cisco MDS 9000 Blade Switch Series, the Cisco MDS 9124, and the Cisco MDS 9134 when NPV mode is enabled.

Configuring Interfaces

- Information About Interfaces, page 2-1
- Prerequisites for Interfaces, page 2-18
- Guidelines and Limitations, page 2-18
- Default Settings, page 2-20
- Configuring Interfaces, page 2-20
- Verifying Interfaces Configuration, page 2-41

Information About Interfaces

The main function of a switch is to relay frames from one data link to another. To relay the frames, the characteristics of the interfaces through which the frames are received and sent must be defined. The configured interfaces can be Fibre Channel interfaces, Gigabit Ethernet interfaces, the management interface (mgmt0), or VSAN interfaces.

This section includes the following topics:

- Interface Description, page 2-2
- Interface Modes, page 2-2
- Interface States, page 2-6
- Graceful Shutdown, page 2-9
- Port Administrative Speeds, page 2-10
- Frame Encapsulation, page 2-10
- Beacon LEDs, page 2-11
- Speed LEDs, page 2-11
- Bit Error Thresholds, page 2-11
- SFP Transmitter Types, page 2-12
- TL Ports, page 2-13
- TL Port ALPA Caches, page 2-14
- Port Guard, page 2-14
- Port Monitor, page 2-15
- Port Monitor Port Guard, page 2-16

- Port Group Monitor, page 2-16
- Local Switching, page 2-16
- Slow Drain Device Detection and Congestion Avoidance, page 2-17
- Management Interfaces, page 2-17
- VSAN Interfaces, page 2-18

Interface Description

For the Fibre Channel interfaces, you can configure the description parameter to provide a recognizable name for the interface. Using a unique name for each interface allows you to quickly identify the interface when you are looking at a listing of multiple interfaces. You can also use the description to identify the traffic or the use for that interface.

Interface Modes

Each physical Fibre Channel interface in a switch may operate in one of several port modes: E port, F port, FL port, TL port, TE port, SD port, ST port, and B port (see Figure 2-1). Besides these modes, each interface may be configured in auto or Fx port modes. These two modes determine the port type during interface initialization.

Figure 2-1 Cisco MDS 9000 Family Switch Port Modes

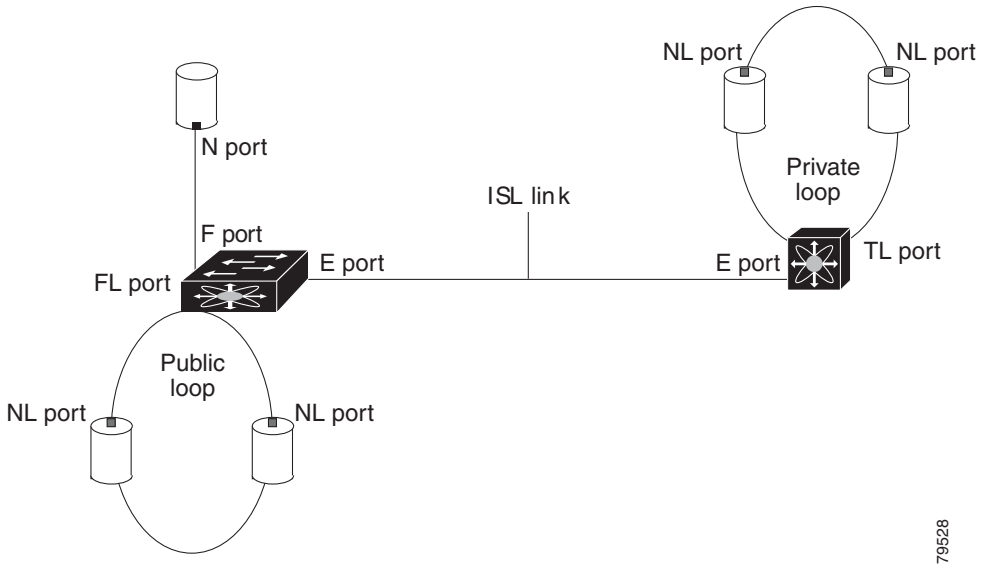

Note Interfaces are created in VSAN 1 by default. See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Each interface has an associated administrative configuration and an operational status:

- The administrative configuration does not change unless you modify it. This configuration has various attributes that you can configure in administrative mode.

- The operational status represents the current status of a specified attribute like the interface speed. This status cannot be changed and is read-only. Some values may not be valid when the interface is down (for example, the operational speed).

Note When a module is removed and replaced with the same type of module, the configuration is retained. If a different type of module is inserted, then the original configuration is no longer retained.

Each interface is briefly described in the sections that follow.

E Port

In expansion port (E port) mode, an interface functions as a fabric expansion port. This port may be connected to another E port to create an Inter-Switch Link (ISL) between two switches. E ports carry frames between switches for configuration and fabric management. They serve as a conduit between switches for frames destined to remote N ports and NL ports. E ports support class 2, class 3, and class F service.

An E port connected to another switch may also be configured to form a PortChannel (see Chapter 6, "Configuring PortChannels").

Note We recommend that you configure E ports on 16-port modules. If you must configure an E port on a 32-port oversubscribed module, then you can only use the first port in a group of four ports (for example, ports 1 through 4, 5 through 8, and so forth). The other three ports cannot be used.

F Port

In fabric port (F port) mode, an interface functions as a fabric port. This port may be connected to a peripheral device (host or disk) operating as an N port. An F port can be attached to only one N port. F ports support class 2 and class 3 service.

FL Port

In fabric loop port (FL port) mode, an interface functions as a fabric loop port. This port may be connected to one or more NL ports (including FL ports in other switches) to form a public arbitrated loop. If more than one FL port is detected on the arbitrated loop during initialization, only one FL port becomes operational and the other FL ports enter nonparticipating mode. FL ports support class 2 and class 3 service.

Note FL port mode is not supported on 4-port 10-Gbps switching module interfaces.

NP Ports

An NP port is a port on a device that is in NPV mode and connected to the core switch via an F port. NP ports function like N ports except that in addition to providing N port operations, they also function as proxies for multiple, physical N ports.

For more details about NP ports and NPV, see Chapter 7, "Configuring N Port Virtualization."

TL Port

In translative loop port (TL port) mode, an interface functions as a translative loop port. It may be connected to one or more private loop devices (NL ports). TL ports are specific to Cisco MDS 9000 Family switches and have similar properties as FL ports. TL ports enable communication between a private loop device and one of the following devices:

- A device attached to any switch on the fabric
- A device on a public loop anywhere in the fabric
- A device on a different private loop anywhere in the fabric
- A device on the same private loop

TL ports support class 2 and class 3 services.

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop (see the "TL Port ALPA Caches" section on page 2-14).

Tip We recommend configuring devices attached to TL ports in zones that have up to 64 zone members.

Note TL port mode is not supported on Generation 2 switching module interfaces.

TE Port

In trunking E port (TE port) mode, an interface functions as a trunking expansion port. It may be connected to another TE port to create an extended ISL (EISL) between two switches. TE ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of E ports to support the following:

- VSAN trunking
- Transport quality of service (QoS) parameters
- Fibre Channel trace (fctrace) feature

In TE port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family switches (see Chapter 5, "Configuring Trunking"). TE ports support class 2, class 3, and class F service.

TF Port

In trunking F port (TF port) mode, an interface functions as a trunking expansion port. It may be connected to another trunked N port (TN port) or trunked NP port (TNP port) to create a link between a core switch and an NPV switch or an HBA to carry tagged frames. TF ports are specific to Cisco MDS 9000 Family switches. They expand the functionality of F ports to support VSAN trunking.

In TF port mode, all frames are transmitted in EISL frame format, which contains VSAN information. Interconnected switches use the VSAN ID to multiplex traffic from one or more VSANs across the same physical link. This feature is referred to as trunking in the Cisco MDS 9000 Family (see Chapter 5, "Configuring Trunking"). TF ports support class 2, class 3, and class F service.

TNP Port

In trunking NP port (TNP port) mode, an interface functions as a trunking expansion port. It may be connected to a trunked F port (TF port) to create a link to a core NPIV switch from an NPV switch to carry tagged frames.

SD Port

In SPAN destination port (SD port) mode, an interface functions as a switched port analyzer (SPAN). The SPAN feature is specific to switches in the Cisco MDS 9000 Family. It monitors network traffic that passes though a Fibre Channel interface. This monitoring is done using a standard Fibre Channel analyzer (or a similar switch probe) that is attached to an SD port. SD ports do not receive frames, they only transmit a copy of the source traffic. The SPAN feature is nonintrusive and does not affect switching of network traffic for any SPAN source ports (see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

ST Port

In the SPAN tunnel port (ST port) mode, an interface functions as an entry point port in the source switch for the RSPAN Fibre Channel tunnel. The ST port mode and the remote SPAN (RSPAN) feature are specific to switches in the Cisco MDS 9000 Family. When configured in ST port mode, the interface cannot be attached to any device, and thus cannot be used for normal Fibre Channel traffic (see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*).

Note ST port mode is not supported on the Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Fx Port

Interfaces configured as Fx ports can operate in either F port or FL port mode. The Fx port mode is determined during interface initialization depending on the attached N port or NL port. This administrative configuration disallows interfaces to operate in any other mode—for example, preventing an interface to connect to another switch.

B Port

While E ports typically interconnect Fibre Channel switches, some SAN extender devices, such as the Cisco PA-FC-1G Fibre Channel port adapter, implement a bridge port (B port) model to connect geographically dispersed fabrics. This model uses B ports as described in the T11 Standard FC-BB-2.

If an FCIP peer is a SAN extender device that only supports Fibre Channel B ports, you need to enable the B port mode for the FCIP link. When a B port is enabled, the E port functionality is also enabled and they coexist. If the B port is disabled, the E port functionality remains enabled (see the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*).

Auto Mode

Interfaces configured in auto mode can operate in one of the following modes: F port, FL port, E port, TE port, or TF port. The port mode is determined during interface initialization. For example, if the interface is connected to a node (host or disk), it operates in F port or FL port mode depending on the N

port or NL port mode. If the interface is attached to a third-party switch, it operates in E port mode. If the interface is attached to another switch in the Cisco MDS 9000 Family, it may become operational in TE port mode (see Chapter 5, "Configuring Trunking").

TL ports and SD ports are not determined during initialization and are administratively configured.

Note Fibre Channel interfaces on Storage Services Modules (SSMs) cannot be configured in auto mode.

Interface States

The interface state depends on the administrative configuration of the interface and the dynamic state of the physical link.

Administrative States

The administrative state refers to the administrative configuration of the interface as described in Table 2-1.

Table 2-1 Administrative States

Administrative State	Description
Up	Interface is enabled.
Down	Interface is disabled. If you administratively disable an interface by shutting down that interface, the physical link layer state change is ignored.

Operational States

The operational state indicates the current operational state of the interface as described in Table 2-2.

Table 2-2 Operational States

Operational State	Description
Up	Interface is transmitting or receiving traffic as desired. To be in this state, an interface must be administratively up, the interface link layer state must be up, and the interface initialization must be completed.
Down	Interface cannot transmit or receive (data) traffic.
Trunking	Interface is operational in TE or TF mode.

Reason Codes

Reason codes are dependent on the operational state of the interface as described in Table 2-3.

Table 2-3 Reason Codes for Interface States

Administrative Configuration	Operational Status	Reason Code
Up	Up	None.
Down	Down	Administratively down—If you administratively configure an interface as down, you disable the interface. No traffic is received or transmitted.
Up	Down	See Table 2-4.

Note Only some of the reason codes are listed in Table 2-4.

If the administrative state is up and the operational state is down, the reason code differs based on the nonoperational reason code as described in Table 2-4.

Table 2-4 Reason Codes for Nonoperational States

Reason Code (long version)	Description	Applicable Modes
Link failure or not connected	The physical layer link is not operational.	All
SFP not present	The small form-factor pluggable (SFP) hardware is not plugged in.	
Initializing	The physical layer link is operational and the protocol initialization is in progress.	
Reconfigure fabric in progress	The fabric is currently being reconfigured.	
Offline	The Cisco NX-OS software waits for the specified R_A_TOV time before retrying initialization.	
Inactive	The interface VSAN is deleted or is in a suspended state. To make the interface operational, assign that port to a configured and active VSAN.	
Hardware failure	A hardware failure is detected.	
Error disabled	Error conditions require administrative attention. Interfaces may be error-disabled for various reasons. For example: • Configuration failure. • Incompatible buffer-to-buffer credit configuration. To make the interface operational, you must first fix the error conditions causing this state; and next, administratively shut down or enable the interface.	
FC redirect failure	A port is isolated because a Fibre Channel redirect is unable to program routes.	
No port activation license available	A port is not active because it does not have a port license.	
SDM failure	A port is isolated because SDM is unable to program routes.	

Table 2-4 Reason Codes for Nonoperational States (continued)

Reason Code (long version)	Description	Applicable Modes
Isolation due to ELP failure	The port negotiation failed.	Only E ports and TE ports
Isolation due to ESC failure	The port negotiation failed.	
Isolation due to domain overlap	The Fibre Channel domains (fcdomain) overlap.	
Isolation due to domain ID assignment failure	The assigned domain ID is not valid.	
Isolation due to the other side of the link E port isolated	The E port at the other end of the link is isolated.	
Isolation due to invalid fabric reconfiguration	The port is isolated due to fabric reconfiguration.	
Isolation due to domain manager disabled	The fcdomain feature is disabled.	
Isolation due to zone merge failure	The zone merge operation failed.	
Isolation due to VSAN mismatch	The VSANs at both ends of an ISL are different.	
Nonparticipating	FL ports cannot participate in loop operations. It may happen if more than one FL port exists in the same loop, in which case all but one FL port in that loop automatically enters nonparticipating mode.	Only FL ports and TL ports
PortChannel administratively down	The interfaces belonging to the PortChannel are down.	Only PortChannel interfaces
Suspended due to incompatible speed	The interfaces belonging to the PortChannel have incompatible speeds.	
Suspended due to incompatible mode	The interfaces belonging to the PortChannel have incompatible modes.	
Suspended due to incompatible remote switch WWN	An improper connection is detected. All interfaces in a PortChannel must be connected to the same pair of switches.	

Graceful Shutdown

Interfaces on a port are shut down by default (unless you modified the initial configuration).

The Cisco NX-OS software implicitly performs a graceful shutdown in response to either of the following actions for interfaces operating in the E port mode:

- If you shut down an interface.
- If a Cisco NX-OS software application executes a port shutdown as part of its function.

A graceful shutdown ensures that no frames are lost when the interface is shutting down. When a shutdown is triggered either by you or the Cisco NX-OS software, the switches connected to the shutdown link coordinate with each other to ensure that all frames in the ports are safely sent through the link before shutting down. This enhancement reduces the chance of frame loss.

A graceful shutdown is not possible in the following situations:

- If you physically remove the port from the switch.
- If in-order delivery (IOD) is enabled (for information about IOD, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*).
- If the Min_LS_interval interval is higher than 10 seconds. For information about FSPF global configuration, refer to the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

Note This feature is only triggered if both switches at either end of this E port interface are MDS switches and are running Cisco SAN-OS Release 2.0(1b) or later, or MDS NX-OS Release 4.1(1a) or later.

Port Administrative Speeds

By default, the port administrative speed for an interface is automatically calculated by the switch.

For internal ports on the Cisco Fabric Switch for HP c_Class BladeSystem and Cisco Fabric Switch for IBM BladeCenter, a port speed of 1 Gbps is not supported. Auto-negotiation is supported between 2 Gbps and 4 Gbps only. Also, if the BladeCenter is a T chassis, then port speeds are fixed at 2 Gbps and auto-negotiation is not enabled.

Autosensing

Autosensing speed is enabled on all 4-Gbps and 8-Gbps switching module interfaces by default. This configuration enables the interfaces to operate at speeds of 1 Gbps, 2 Gbps, or 4 Gbps on the 4-Gbps switching modules, and 8 Gbps on the 8-Gbps switching modules. When autosensing is enabled for an interface operating in dedicated rate mode, 4 Gbps of bandwidth is reserved, even if the port negotiates at an operating speed of 1 Gbps or 2 Gbps.

To avoid wasting unused bandwidth on 48-port and 24-port 4-Gbps and 8-Gbps Fibre Channel switching modules, you can specify that only 2 Gbps of required bandwidth be reserved, not the default of 4 Gbps or 8 Gbps. This feature shares the unused bandwidth within the port group provided that it does not exceed the rate limit configuration for the port. You can also use this feature for shared rate ports that are configured for autosensing.

Tip When migrating a host that supports up to 2-Gbps traffic (that is, not 4 Gbps with autosensing capabilities) to the 4-Gbps switching modules, use autosensing with a maximum bandwidth of 2 Gbps. When migrating a host that supports up to 4-Gbps traffic (that is, not 8 Gbps with autosensing capabilities) to the 8-Gbps switching modules, use autosensing with a maximum bandwidth of 4 Gbps.

Frame Encapsulation

The **switchport encap eisl** command only applies to SD port interfaces. This command determines the frame format for all frames transmitted by the interface in SD port mode. If the encapsulation is set to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources.

The **switchport encap eisl** command is disabled by default. If you enable encapsulation, all outgoing frames are encapsulated, and you will see a new line (Encapsulation is eisl) in the **show interface** *SD_port_interface* command output. See the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

You can set the frame format to EISL for all frames transmitted by the interface in SD port mode. If you sent the frame encapsulation to EISL, all outgoing frames are transmitted in the EISL frame format, regardless of the SPAN sources. See the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

Beacon LEDs

Figure 2-2 displays the status, link, and speed LEDs in a 16-port switching module.

Figure 2-2 Cisco MDS 9000 Family Switch Interface Modes

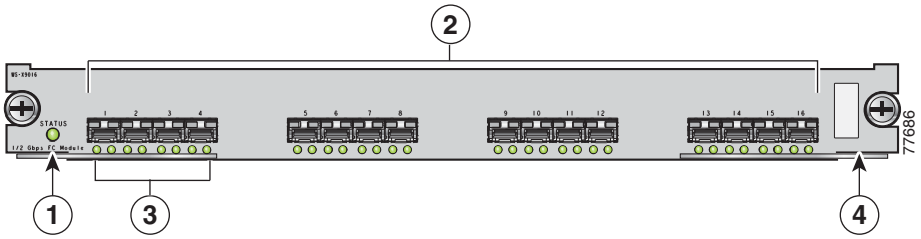

1	Status LED[1]	**3**	Link LEDs[1] and speed LEDs[2]
2	1/2-Gbps Fibre Channel port group[3]	**4**	Asset tag[4]

1. See the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.
2. See the "Speed LEDs" section on page 2-11.
3. See the "Generation 1 Interface Configuration Guidelines" section on page 2-18.
4. Refer to the Cisco MDS 9000 Family hardware installation guide for your platform.

Speed LEDs

Each port has one link LED on the left and one speed LED on the right.

The speed LED displays the speed of the port interface:

- Off—The interface attached to that port is functioning at 1000 Mbps.
- On (solid green)—The interface attached to that port is functioning at 2000 Mbps (for 2 Gbps interfaces).

The speed LED also displays if the beacon mode is enabled or disabled:

- Off or solid green—Beacon mode is disabled.
- Flashing green—The beacon mode is enabled. The LED flashes at one-second intervals.

Note Generation 2, Generation 3, and Generation 4 modules and fabric switches do not have speed LEDs.

Bit Error Thresholds

The bit error rate threshold is used by the switch to detect an increased error rate before performance degradation seriously affects traffic.

The bit errors can occur for the following reasons:

- Faulty or bad cable.
- Faulty or bad GBIC or SFP.
- GBIC or SFP is specified to operate at 1 Gbps but is used at 2 Gbps.
- GBIC or SFP is specified to operate at 2 Gbps but is used at 4 Gbps.
- Short haul cable is used for long haul or long haul cable is used for short haul.
- Momentary sync loss.
- Loose cable connection at one or both ends.
- Improper GBIC or SFP connection at one or both ends.

A bit error rate threshold is detected when 15 error bursts occur in a 5-minute period. By default, the switch disables the interface when the threshold is reached. You can enter a **shutdown** and **no shutdown** command sequence to re-enable the interface.

You can configure the switch to not disable an interface when the threshold is crossed. By default, the threshold disables the interface.

SFP Transmitter Types

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed. Table 2-5 defines the acronyms used for SFPs.

The small form-factor pluggable (SFP) hardware transmitters are identified by their acronyms when displayed in the **show interface brief** command. If the related SFP has a Cisco-assigned extended ID, then the **show interface** and **show interface brief** commands display the ID instead of the transmitter type. The **show interface transceiver** command and the **show interface fc** *slot*/*port* **transceiver** command display both values for Cisco-supported SFPs. Table 2-5 defines the acronyms used in the command output (see the "Displaying Interface Information" section on page 2-42).

Table 2-5 SFP Transmitter Acronym Definitions

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
short wave laser	swl
long wave laser	lwl
long wave laser cost reduced	lwcr
electrical	elec
Extended transmitters assigned to Cisco-supported SFPs	
CWDM-1470	c1470
CWDM-1490	c1490
CWDM-1510	c1510
CWDM-1530	c1530
CWDM-1550	c1550
CWDM-1570	c1570

Table 2-5 SFP Transmitter Acronym Definitions (continued)

Definition	Acronym
Standard transmitters defined in the GBIC specifications	
CWDM-1590	c1590
CWDM-1610	c1610

TL Ports

Private loop devices refer to legacy devices that reside on arbitrated loops. These devices are not aware of a switch fabric because they only communicate with devices on the same physical loop. The legacy devices are used in Fibre Channel networks, and devices outside the loop may need to communicate with them. The communication functionality is provided through TL ports. See the "Interface Modes" section on page 2-2.

TL port mode is not supported on the following hardware:

- Generation 2 switching module interfaces
- Cisco MDS 9124 Fabric Switch
- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter

Table 2-6 lists the TL port translations supported in Cisco MDS 9000 Family switches.

Table 2-6 Supported TL Port Translations

Translation from	Translation to	Example
Private initiator	Private target	From I1 to T1 or vice versa
Private initiator	Public target — N port	From I1 to T2 or vice versa
Private initiator	Public target — NL port	From I4 to T3 or vice versa
Public initiator — N port	Private target	From I2 to T1 or vice versa
Public initiator — NL port	Private target	From I3 to T1 or vice versa

Figure 2-3 shows examples of TL port translation support.

Figure 2-3 TL Port Translation Support Examples

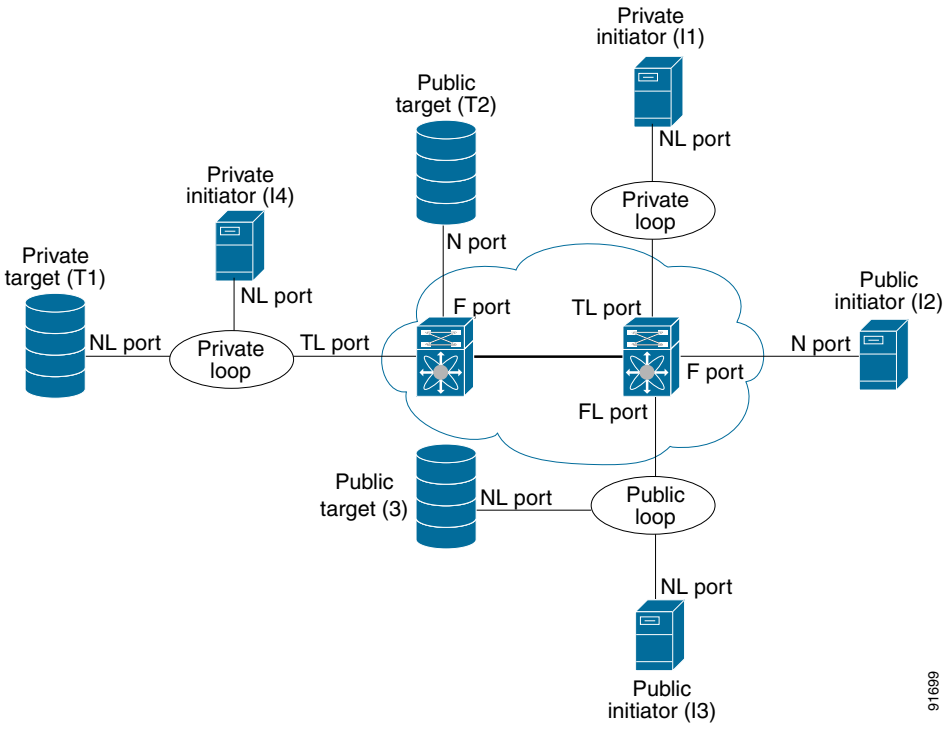

TL Port ALPA Caches

Although TL ports cannot be automatically configured, you can manually configure entries in arbitrated loop physical address (ALPA) caches. Generally, ALPA cache entries are automatically populated when an ALPA is assigned to a device. Each device is identified by its port world wide name (pWWN). When a device is allocated an ALPA, an entry for that device is automatically created in the ALPA cache.

A cache contains entries for recently allocated ALPA values. These caches are maintained on various TL ports. If a device already has an ALPA, the Cisco NX-OS software attempts to allocate the same ALPA to the device each time. The ALPA cache is maintained in persistent storage and saves information across switch reboots. The maximum cache size is 1000 entries. If the cache is full, and a new ALPA is allocated, the Cisco NX-OS software discards an inactive cache entry (if available) to make space for the new entry. See the "TL Port" section on page 2-4 for more information on TL ports.

Port Guard

The port guard feature is intended for use in environments where the system and application environment does not adapt quickly and efficiently to a port going down and back up, or to a port rapidly cycling up and down, which can happen in some failure modes. For example, if a system takes five seconds to stabilize after a port goes down, but the port is going up and down once a second, a more severe failure in the fabric might occur.

The port guard feature gives the SAN administrator the ability to prevent this issue from occurring in environments that are vulnerable to these problems. The port can be configured to stay down after the first failure or after a specified number of failures in a specified time period. This allows the SAN administrator to intervene and control the recovery, avoiding any problems caused by the cycling.

Using the port guard feature, you can restrict the number of error reports and bring a malfunctioning port to down state dynamically. A port can be configured to go into error-disabled state for specific types of failures.

A general link failure caused by link-down is the superset of all other causes. The sum of the number of all other causes equals to the number of link-down link failures. This means a port is brought to down state when it reaches the maximum number of allowed link failures or the number of specific causes.

The causes of link failure can be any of the following:

- ESP trustsec-violation
- Bit-errors
- Signal loss
- Sync loss
- Link reset
- Credit loss
- Additional causes might be the following:
  - Not operational (NOS).
  - Too many interrupts.
  - Cable is disconnected.
  - Hardware recoverable errors.
  - The connected device rebooted (F ports only).
  - The connected linecard rebooted (ISL only).

Port Monitor

Port monitor helps to monitor the performance and the status of ports and generate alerts when problems occur. You can configure the thresholds for various counters and trigger an event when the values cross the threshold settings.

The default port monitor policy has the following threshold values:

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	Event	% Falling Threshold	Event
Link Loss	Delta	60	5	4	1	4
Sync Loss	Delta	60	5	4	1	4
Protocol Error	Delta	60	1	4	0	4
Signal Loss	Delta	60	5	4	1	4
Invalid Words	Delta	60	1	4	0	4
Invalid CRCs	Delta	60	5	4	1	4

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	Event	% Falling Threshold	Event
RX Performance	Delta	60	2147483648	4	524288000	4
TX Performance	Delta	60	2147483648	4	524288000	4

Port Monitor Port Guard

Port monitor port guard is a feature that disables or shuts down a port when an event occurs. Depending on the configuration, when an event occurs the port is either error-disabled or flapped.

Port monitor port guard is a different or separate feature that functions based on the configuration of the **errordisable** command.

Port Group Monitor

Each line card or module has a predefined set of ports which share the same backplane bandwidth called port groups. While oversubscription is a feature, the port group monitor feature helps to monitor the spine bandwidth utilization. An alarm syslog is generated so that you can provision the ports across port groups evenly to manage the oversubscription better.

When the port group monitor feature is enabled and a policy consisting of polling interval in seconds, and the raising and falling thresholds in percentage are specified, port group monitor generates a syslog if a port group traffic goes above the specified percentage of the maximum supported bandwidth for that port group (for rx and for tx) and another syslog if the value falls below the specified threshold.

The default port group policy has the following threshold values:

Counter	Threshold Type	Interval (Seconds)	% Rising Threshold	% Falling Threshold
RX Performance	Delta	60	80	20
TX Performance	Delta	60	80	20

Local Switching

Local switching can be enabled in Generation 4 modules, which allows traffic to be switched directly with a local crossbar when the traffic is directed from one port to another on the same line card. By using local switching, an extra switching step is avoided, which decreases the latency.

When using local switching, note the following guidelines:

- All ports need to be in shared mode, which usually is the default state. To place a port in shared mode, enter the **switchport ratemode shared** command.
- E ports are not allowed in the module because they must be in dedicated mode.

Note Local Switching is not supported on MDS 9710.

Slow Drain Device Detection and Congestion Avoidance

All data traffic between end devices in a SAN fabric is carried by Fibre Channel Class 3. In some cases, the traffic is carried by Class 2 services that use link-level, per-hop-based, and buffer-to-buffer flow control. These classes of service do not support end-to-end flow control. When there are slow devices attached to the fabric, the end devices do not accept the frames at the configured or negotiated rate. The slow devices lead to ISL credit shortage in the traffic destined for these devices and they congest the links. The credit shortage affects the unrelated flows in the fabric that use the same ISL link even though destination devices do not experience slow drain.

This feature provides various enhancements to detect slow drain devices that are causing congestion in the network and also provides a congestion avoidance function.

This feature is focused mainly on the edge ports that are connected to slow drain devices. The goal is to avoid or minimize the frames being stuck in the edge ports due to slow drain devices that are causing ISL blockage. To avoid or minimize the stuck condition, configure lesser frame timeout for the ports. No-credit timeout drops all packets once the slow drain is detected using the configured thresholds. The lesser frame timeout value helps to alleviate the slow drain condition that affects the fabric by dropping the packets on the edge ports sooner than the time they actually get timed out (500 ms). This function frees the buffer space in ISL, which can be used by other unrelated flows that do not experience slow drain condition.

Note This feature is used mainly for edge ports that are connected to slow edge devices. Even though this feature can be applied to ISLs as well, we recommend that you apply this feature only for edge F ports and retain the default configuration for ISLs as E and TE ports. This feature is not supported on Generation 1 modules.

Management Interfaces

You can remotely configure the switch through the management interface (mgmt0). To configure a connection on the mgmt0 interface, you must configure either the IP version 4 (IPv4) parameters (IP address, subnet mask, and default gateway) or the IP version 6 (IPv6) parameters so that the switch is reachable.

Before you begin to configure the management interface manually, obtain the switch's IPv4 address and subnet mask, or the IPv6 address.

The management port (mgmt0) is autosensing and operates in full-duplex mode at a speed of 10/100/1000 Mbps. Autosensing supports both the speed and the duplex mode. On a Supervisor-1 module, the default speed is 100 Mbps and the default duplex mode is auto. On a Supervisor-2 module, the default speed is auto and the default duplex mode is auto.

Note You need to explicitly configure a default gateway to connect to the switch and send IP packets or add a route for each subnet.

VSAN Interfaces

VSANs apply to Fibre Channel fabrics and enable you to configure multiple isolated SAN topologies within the same physical infrastructure. You can create an IP interface on top of a VSAN and then use this interface to send frames to this VSAN. To use this feature, you must configure the IP address for this VSAN. VSAN interfaces cannot be created for nonexisting VSANs.

Prerequisites for Interfaces

Before you begin configuring the interfaces, ensure that the modules in the chassis are functioning as designed. To verify the status of a module at any time, enter the **show module** command in EXEC mode. For information about verifying the module status, refer to the *Cisco NX-OS Fundamentals Configuration Guide*.

Guidelines and Limitations

This section includes the following topics:

- Generation 1 Interface Configuration Guidelines, page 2-18
- Private Loop Configuration Guidelines, page 2-19
- VSAN Interface Configuration Guidelines, page 2-19

Generation 1 Interface Configuration Guidelines

The Generation 1 interfaces configuration guidelines apply to the following hardware:

- The 32-port, 2-Gbps or 1-Gbps switching module interfaces
- The Cisco MDS 9140 and 9120 switch interfaces

Note Due to the hardware design of the MDS 9134 switch, we do not support interface out-of-service action on either of its two 10-Gigabit ports. This is because no internal port hardware resource is released when an out-of-service action is performed on these 10-Gigabit ports.

When configuring these host-optimized ports, the following port mode guidelines apply:

- You can configure only the first port in each 4-port group (for example, the first port in ports 1-4, the fifth port in ports 5-8, and so on) as an E port. If the first port in the group is configured as an E port, the other three ports in each group (ports 2-4, 6-8, and so on) are not usable and remain shutdown.
- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- If any of the other three ports are enabled, you cannot configure the first port as an E port. The other three ports continue to remain enabled.

- The auto mode is not allowed in a 32-port switching module or the host-optimized ports in the Cisco 9100 Series (16 host-optimized ports in the Cisco MDS 9120 switch and 32 host-optimized ports in the Cisco MDS 9140 switch).
- The default port mode is Fx (Fx negotiates to F or FL) for 32-port switching modules.
- The 32-port switching module does not support FICON.

Note We recommend that you configure your E ports on a 16-port switching module. If you must configure an E port on a 32-port host-optimized switching module, the other three ports in that 4-port group cannot be used.

Note In the Cisco MDS 9100 Series, the groups of ports that are located on the left and outlined in white are full line rate. The other ports are host-optimized. Each group of 4 host-optimized ports have the same features as for the 32-port switching module.

Private Loop Configuration Guidelines

Follow these guidelines when configuring private loops:

- A maximum of 64 fabric devices can be proxy to a private loop.
- Fabric devices must be in the same zone as private loop devices to be proxy to the private loop.
- Each private device on a TL port may be included in a different zone.
- All devices on the loop are treated as private loops. You cannot mix private and public devices on the loop if the configured port mode is TL.
- The only FC4-type supported by TL ports is SCSI (FCP).
- Communication between a private initiator to a private target on the same private loop does not invoke TL port services.

VSAN Interface Configuration Guidelines

Follow these guidelines when creating or deleting VSAN interfaces:

- Create a VSAN before creating the interface for that VSAN. If a VSAN does not exist, the interface cannot be created.
- Create the interface VSAN—it is not created automatically.
- If you delete the VSAN, the attached interface is automatically deleted.
- Configure each interface only in one VSAN.

Tip After configuring the VSAN interface, you can configure an IP address or Virtual Router Redundancy Protocol (VRRP) feature. See the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

Default Settings

Table 2-7 lists the default settings for interface parameters.

Table 2-7 Default Interface Parameters

Parameters	Default
Interface mode	Auto
Interface speed	Auto
Administrative state	Shutdown (unless changed during initial setup)
Trunk mode	On (unless changed during initial setup) on non-NPV and NPIV core switches. Off on NPV switches.
Trunk-allowed VSANs or VF-IDs	1 to 4093
Interface VSAN	Default VSAN (1)
Beacon mode	Off (disabled)
EISL encapsulation	Disabled
Data field size	2112 bytes

Configuring Interfaces

This section includes the following topics:

- Configuring Fibre Channel Interfaces, page 2-21
- Setting the Interface Administrative State, page 2-22
- Configuring Interface Modes, page 2-22
- Configuring System Default Port Mode F, page 2-23
- Configuring Port Administrative Speeds, page 2-25
- Configuring Port Speed Group, page 2-26
- Configuring the Interface Description, page 2-26
- Specifying a Port Owner, page 2-27
- Configuring Beacon Mode, page 2-27
- Disabling Bit Error Threshold, page 2-28
- Configuring Switch Port Attribute Default Values, page 2-28
- Configuring TL Ports, page 2-29
- Manually Inserting Entries into the ALPA Cache, page 2-29
- Clearing the ALPA Cache, page 2-29
- Configuring Port Guard, page 2-30
- Configuring Port Monitor, page 2-31
- Configuring Port Group Monitor, page 2-34

- Configuring Management Interfaces, page 2-37
- Creating VSAN Interfaces, page 2-38
- Configuring Slow Drain Device Detection and Congestion Avoidance, page 2-38

For more information on configuring mgmt0 interfaces, refer to the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide* and *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

For more information on configuring Gigabit Ethernet interfaces, see the *Cisco MDS 9000 Family NX-OS IP Services Configuration Guide*.

Configuring Fibre Channel Interfaces

Detailed Steps

To configure a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode. **Note** When a Fibre Channel interface is configured, it is automatically assigned a unique world wide name (WWN). If the interface's operational state is up, it is also assigned a Fibre Channel ID (FC ID).

To configure a range of interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1 - 4 , fc2/1 - 3` `switch(config-if)#`	Selects the range of Fibre Channel interfaces and enters interface configuration submode3. **Note** In this command, provide a space before and after the comma.

For the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter, you can configure a range of interfaces among internal ports or external ports, but you cannot mix both interface types within the same range. For example, "bay 1-10, bay 12" or "ext 0, ext 15-18" are valid ranges, but "bay 1-5, ext 15-17" is not.

Setting the Interface Administrative State

Detailed Steps

To gracefully shut down an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# shutdown`	Gracefully shuts down the interface and administratively disables traffic flow (default).

To enable traffic flow, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# no shutdown`	Enables traffic flow to administratively allow traffic when the **no** prefix is used (provided the operational state is up).

Configuring Interface Modes

Detailed Steps

To configure the interface mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport mode F` `switch(config-if)#`	Configures the administrative mode of the port. You can set the operational state to auto, E, F, FL, Fx, TL, NP, or SD port mode. **Note** Fx ports refers to an F port or an FL port (host connection only), but not E ports.
	`switch(config-if)# switchport mode auto` `switch(config-if)#`	Configures the interface mode to auto-negotiate an E, F, FL, or TE port mode (not TL or SD port modes) of operation. **Note** TL ports and SD ports cannot be configured automatically. They must be administratively configured. **Note** You cannot configure Fibre Channel interfaces on SSMs in auto mode.

Configuring System Default Port Mode F

The **system default switchport mode F** command sets the administrative mode of all Fibre Channel ports to mode F, while avoiding traffic disruption caused by the formation of unwanted Inter-Switch Links (ISLs). This command is part of the setup utility that runs during bootup after a **write erase** or **reload**. It can also be executed from the command line in configuration mode. This command changes the configuration of the following ports to administrative mode F:

- All ports that are down and that are not out-of-service.
- All F ports that are up, whose operational mode is F, and whose administrative mode is not F.

This command does not affect the configuration of the following ports:

- All user-configured ports, even if they are down.
- All non-F ports that are up; however, if non-F ports are down, this command changes the administrative mode of those ports.

Restrictions

- To ensure that ports that are part of ISLs do not get changed to port mode F, configure the ports in port mode E, rather than in Auto mode.
- When the command is executed from the command line, switch operation remains graceful. No ports are flapped.

Detailed Steps

To set the administrative mode of Fibre Channel ports to mode F in the CLI, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system default switchport mode F`	Sets the administrative mode of Fibre Channel ports to mode F (if applicable).
	`switch(config)# no system default switchport mode F`	Sets the administrative mode of Fibre Channel ports to the default (unless user configured).

Note For detailed information about the switch setup utility, see the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.

Example 2-1 shows the command in the setup utility, and Example 2-2 shows the command from the command line.

Example 2-1 Setup Utility

```
Configure default switchport mode F (yes/no) [n]: y
```

Example 2-2 Command Line

```
switch(config)# system default switchport mode F
```

Configuring ISL between Two Switches

Note Ensure the FC cable is connected between the ports and do a no-shut on each port.

E-port mode is used when the port functions as one end of an Inter-Switch Link (ISL) setting. When you set the port mode to E, you restrict the port coming up as an E port (trunking or non-trunking, depending on the trunking port mode).

```
switch# conf t --------> Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# interface fc <Slot No/Port No>
switch(config-if)# switchport mode E
switch(config-if)# end
```

Ensure you follow this on both the switches between which you are attempting to bring up the ISL link.

Configuring 10-Gbps FC Mode

The 48-port 8-Gbps Advanced Fibre Channel module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9)can switch between two speed modes—the 1-, 2-, 4-, 8-Gbps or 10-Gbps. By default, the modules are online in the 1-, 2-, 4-, and 8-Gbps modes when they are loaded for the first time. There are two ways to change the ports to the 10-Gbps speed mode:

- Using the **10G-speed mode** command, which is the recommended method.
- Using the generic speed configuration **switchport speed** command which has certain constraints.

The following conditions apply when the ports in the module can be configured to 10-Gbps speed mode:

- The ports in the module can be configured to 10-Gbps speed only when the DS-13SLT-FAB3 module bandwidth is 256-G. Any other combination of fabric modules will not let the ports come up in 10-Gbps.
- When in 10-Gbps mode, the ports in the module that are not 10-Gbps capable are disabled and will be in out-of-service state.
- The ports function only in full rate mode. They cannot be moved to shared rate mode.
- The ports cannot be configured in any other speed.
- Ports that are capable of 10-Gbps that are disabled or out-of-service cannot be put back in service using the **no out-of-service** command. To put these ports back in service, all ports in the module first have to be moved to the out-of-service state. Then they can be brought back to the in service state.
- Local switching must be disabled, otherwise, ports cannot be configured in dedicated mode.

Only certain ports on the 48-port and 32-port 8-Gbps Advanced Fibre Channel modules are 10-Gbps capable. When running in 10-Gbps mode, the non-10-Gbps ports cannot be operational. They have to be either in shut state or out-of-service state.

Detailed Steps

To configure the interface mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1-12` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode. **Note** Ensure that one full ASIC range of ports are selected before executing this command. For example, /1-12 for a 48-port module or fcY/1-8 for a 32-port module. The first ASIC on the 48-port module has ports fc1/1-12.
Step 3	`switch(config-if)# 10G-speed-mode`	Configures all the ports (fc1 to 12) to out-of-service state. Moves the ports that are capable of a10-Gbps configuration (fc1/4-8 and fc1/10) to in-service state. Sets the speed on ports fc1/4-8 and fc1/10 to 10-Gbps. Sets port modes on these ports to dedicated. Performs a no shut on these ports.
	`switch(config-if)# no 10G-speed-mode`	Reverts the settings and puts all the ports (fc1 to 12) in out-of-service state and moves them to in-service state.

Configuring Port Administrative Speeds

Restrictions

- Changing the port administrative speed is a disruptive operation.

Detailed Steps

To configure the port speed of the interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1`	Selects the fc interface and enters interface configuration mode.

	Command	Purpose
Step 3	`switch(config-if)# switchport speed 1000`	Configures the port speed of the interface to 1000 Mbps. The number indicates the speed in megabits per second (Mbps).
	`switch(config-if)# switchport speed 10000`	Configures the port speed of the interface to 10000 Mbps (for 10-Gbps). **Note** All the 10-Gbps capable interfaces, except the interface that is being configured, must be in the out-of-service state. At least one other 10-Gbps capable interface must be in the in-service state.
	`switch(config-if)# no switchport speed`	Reverts the factory default (auto) administrative speed of the interface.

Configuring Port Speed Group

Detailed Steps

To configure the port speed group of the interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1`	Selects the fc interface and enters interface configuration mode.
Step 3	`switch(config-if)# speed group 10g`	Configures the port speed group to 10 Gbps. **Note** The preferred way of changing the speed group is the **10g-speed-mode** command.
	`switch(config-if)# no speed group 10g`	Unsets the port speed group and reverts to the factory default (auto) administrative speed group of the interface.

Configuring the Interface Description

The interface description can be any alphanumeric string.

Detailed Steps

To configure a description for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport description cisco-HBA2`	Configures the description of the interface. The string can be up to 80 characters long.
	`switch(config-if)# no switchport description`	Clears the description of the interface.

Specifying a Port Owner

Using the port owner feature, you can specify the owner of a port and the purpose for which a port is used so that the other administrators are informed.

Note The port guard and port owner features are available for all ports regardless of the operational mode.

Detailed Steps

To specify or remove the port owner, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects the port interface.
Step 3	`switch(config)# switchport owner` *description*	Specifies the owner of the switch port. The description can include name of the owner and the purpose for which the port is used. The description can be up to 80 characters long.
	`switch(config)# no switchport owner`	Removes (default) the port owner description.

To display the owner description specified for a port, use the following commands:

```
switch# show running interface fc module-number/interface-number
switch# show port internal info interface fc module-number/interface-number
```

Configuring Beacon Mode

By default, the beacon mode is disabled on all switches. The beacon mode is indicated by a flashing green light that helps you identify the physical location of the specified interface. Configuring the beacon mode has no effect on the operation of the interface.

Detailed Steps

To enable beacon mode for a specified interface or range of interfaces, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport beacon`	Enables the beacon mode for the interface.
	`switch(config-if)# no switchport beacon`	Disables the beacon mode for the interface.

Troubleshooting Tips

- The flashing green light turns on automatically when an external loopback is detected that causes the interfaces to be isolated. The flashing green light overrides the beacon mode configuration. The state of the LED is restored to reflect the beacon mode configuration after the external loopback is removed.

Disabling Bit Error Threshold

Detailed Steps

To disable the bit error threshold for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport ignore bit-errors`	Prevents the detection of bit error threshold events from disabling the interface.
	`switch(config-if)# no switchport ignore bit-errors`	Prevents the detection of bit error threshold events from enabling the interface.

Troubleshooting Tips

- Regardless of the setting of the **switchport ignore bit-errors** command, the switch generates a syslog message when bit-error threshold events are detected.

Configuring Switch Port Attribute Default Values

You can configure attribute default values for various switch port attributes. These attributes will be applied globally to all future switch port configurations, even if you do not individually specify them at that time.

Detailed Steps

To configure switch port attributes, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# no system default switchport shutdown` `switch(config)#`	Configures the default setting for administrative state of an interface as Up. (The factory default setting is Down). **Tip** This command is applicable only to interfaces for which no user configuration exists for the administrative state.
	`switch(config)# system default switchport shutdown` `switch(config)#`	Configures the default setting for administrative state of an interface as Down. This is the factory default setting. **Tip** This command is applicable only to interfaces for which no user configuration exists for the administrative state.
	`switch(config)# system default switchport trunk mode auto` `switch(config)#`	Configures the default setting for administrative trunk mode state of an interface as Auto. **Note** The default setting is trunk mode on.

Configuring TL Ports

Private loops require setting the interface mode to TL.

Use the **switchport mode** command to configure a TL port. See the "Configuring Interface Modes" section on page 2-22.

Manually Inserting Entries into the ALPA Cache

Detailed Steps

To manually insert entries into the ALPA cache, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# tlport alpa-cache interface fc1/2 pwwn 22:00:00:20:37:46:09:bd alpa 0x02`	Configures manual entries into the ALPA cache.
Step 3	`switch(config)# tlport alpa-cache interface fc1/3 pwwn 22:00:00:20:37:46:09:bd`	Removes this entry from the ALPA cache.

Clearing the ALPA Cache

The **clear tlport alpa-cache** command clears the entire content of the ALPA cache.

Configuring Port Guard

Detailed Steps

To enable or disable the port guard for a port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1`	Selects the port interface.
Step 3	`switch(config)# errdisable detect cause link-down`	Brings the port to down state if the link flaps once.
	`switch(config)# errdisable detect cause link-down [num-times` *number* `duration` *seconds*`]`	Enables the port guard configuration for the interface. Brings the port to down state if the link flaps for the *number* of instances within the specified *seconds*.
	`switch(config)# no errdisable detect cause link-down`	Removes (default) the port guard configuration for the interface. The link resumes flapping and sending error reports normally.
Step 4	`switch(config)# errdisable detect cause {trustsec-violation | bit-errors | credit-loss | link-reset | signal-loss | sync-loss}`	Brings the port to down state if the specified error occurs even once.
	`switch(config)# errdisable detect cause {trustsec-violation | bit-errors | credit-loss | link-reset | signal-loss | sync-loss} [num-times` *number* `duration` *seconds*`]`	Brings the port to down state if the specified error occurs for the *number* of instances within the specified *seconds*.
	`switch(config)# no errdisable detect cause {trustsec-violation | bit-errors | credit-loss | link-reset | signal-loss | sync-loss}`	Removes (default) the port guard configuration for the interface. The link resumes flapping and sending error reports normally.

Examples

This example shows how to configure port guard to bring a port to down state if the link flaps 5 times within 120 seconds based on multiple causes:

```
Switch# config t
Switch (config)# interface fc1/1
Switch (config-if)# errdisable detect cause link-down num-times 5 duration 120
Switch (config-if)# errdisable detect cause bit-errors num-times 5 duration 120
Switch (config-if)# errdisable detect cause credit-loss num-times 5 duration 120
```

The above example sets the configuration to the following status:

- The port will be error-disabled due to bit errors if the port suffers link failure due to bit errors 5 times in 120 seconds.
- The port will be error-disabled due to credit loss if the port suffers link failure due to credit loss 5 times in 120 seconds.
- The port will be error-disabled due to link down if the port suffers link failure due to bit errors 2 times and link-failure due to credit loss 3 times in 120 seconds.

This example shows the internal information about a port in down state because of TrustSec violation:

```
Switch# show port internal info interface fc8/3
fc8/3 is down (Error disabled - port down due to trustsec violation)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 21:c3:00:0d:ec:10:57:80
    Admin port mode is E, trunk mode is on
    snmp link state traps are enabled
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      11274 frames input, 1050732 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      11242 frames output, 971900 bytes
        0 discards, 0 errors
      11 input OLS, 34 LRR, 10 NOS, 0 loop inits
      72 output OLS, 37 LRR, 2 NOS, 0 loop inits
    Interface last changed at Sun Nov 27 07:34:05 1988

   admin port-down trustsec-violation(3) num_times 0, duration = 0
    state reason (Error disabled - port down due to trustsec violation)
    Port guard trustsec violation is Enabled
    errdisabled on trustsec violation TRUE, oper cnt = 1
    port guard first trustsec violation Sun Nov 27 07:34:05 1988
```

Troubleshooting Tips

- Link down is the superset of all other causes. A port is brought to down state if the total number of other causes equals to the number of allowed link-down failures.
- Even if the link does not flap due to failure of the link, and port guard is not enabled, the port goes into a down state if too many invalid FLOGI requests are received from the same host. Use the **shut** and the **no shut** commands consecutively to bring up the link.

Configuring Port Monitor

This section includes the following topics:

- Enabling Port Monitor, page 2-31
- Configuring a Port Monitor Policy, page 2-32
- Activating a Port Monitor Policy, page 2-34

Enabling Port Monitor

Detailed Steps

To enable port monitor, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-monitor enable`	Enables (default) port monitoring.
	`switch(config)# no port-monitor enable`	Disables port monitoring.

Configuring a Port Monitor Policy

Detailed Steps

To configure a port monitor policy, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-monitor name` *policyname*	Specifies the policy name and enters the port monitoring policy configuration mode.
	`switch(config)# no port-monitor name` *policyname*	Removes the policy.
	`switch# config t`	Enters configuration mode.
Step 3	`switch(config-port-monitor)# port-type access-port`	Applies the policy to the access ports.
	`switch(config-port-monitor)# port-type trunks`	Applies the policy to the trunk ports.
	`switch(config-port-monitor)# port-type all`	Applies the policy to all ports.
Step 4	`switch(config-port-monitor)# counter credit-loss-reco poll-interval` *seconds* *percentage1* `event` *event-id* `delta rising-threshold` *percentage1* `event` *event-id*	Specifies the delta credit loss recovery counter, poll interval in seconds, the thresholds in percentage, and the event IDs of events.
Step 5	`switch(config-port-monitor)# monitor counter err-pkt-from-port`[1] `switch(config-port-monitor)# counter err-pkt-from-port`[1] `poll-interval` *seconds* `delta rising-threshold` *eventcount1* `event` *event-id* `falling-threshold`[2] *eventcount2* `event` *event-id*	Activates the err-pkt-from-port counter. Specifies the delta err-pkt-from-port counter, poll interval in seconds, the thresholds in percentage, and the event IDs of events.
Step 6	`switch(config-port-monitor)# monitor counter err-pkt-from-xbar`[1] `switch(config-port-monitor)# counter err-pkt-from-xbar`[1] `poll-interval` *seconds* `delta rising-threshold` *eventcount1* `event` *event-id*	Activates the err-pkt-from-xbar counter. Specifies the delta err-pkt-from-xbar counter, poll interval in seconds, the thresholds in the count of error frame events, and the event IDs of events.
Step 7	`switch(config-port-monitor)# monitor counter err-pkt-to-xbar`[1] `switch(config-port-monitor)# counter err-pkt-to-xbar`[1] `poll-interval` *seconds* `delta rising-threshold` *eventcount1* `event` *event-id* `falling-threshold`[2] *eventcount2* `event` *event-id*	Activates the err-pkt-to-xbar counter. Specifies the delta err-pkt-to-xbar counter, poll interval in seconds, the thresholds in the count of error frame events, and the event IDs of events.
Step 8	`switch(config-port-monitor)# counter invalid-crc poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta invalid CRC, poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 9	`switch(config-port-monitor)# counter invalid-words poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta invalid words, poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.

	Command	Purpose
Step 10	`switch(config-port-monitor)# counter link-loss poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta link failure counter, poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 11	`switch(config-port-monitor)# counter protocol-error poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta protocol error poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 12	`switch(config-port-monitor)# counter rx-performance poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta Rx counter poll interval in seconds and thresholds in percentage.
Step 13	`switch(config-port-monitor)# counter signal-loss poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta signal loss poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 14	`switch(config-port-monitor)# counter state-change poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta state change poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 15	`switch(config-port-monitor)# counter sync-loss poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta sync loss poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered.
Step 16	`switch(config-port-monitor)# counter tx-performance poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold`[2] *percentage2* `event` *event-id*	Specifies the delta Tx counter poll interval in seconds and thresholds in percentage.
Step 17	`switch(config-port-monitor)# no counter sync-loss`	Reverts to the default policy for sync loss performance counter values.
Step 18	`switch(config-port-monitor)# no counter tx-performance`	Reverts to the default policy for Tx performance counter values.
Step 19	`switch(config-port-monitor)# monitor counter rx-datarate`	Turns on Rx performance counter.
	`switch(config-port-monitor)# monitor counter tx-datarate`	Turns on Tx performance counter.
	`switch(config-port-monitor)# no monitor counter tx-datarate`	Turns off Tx performance counter.
	`switch(config-port-monitor)# no monitor counter sync-loss`	Turns off monitoring sync loss.
	`switch(config-port-monitor)# no monitor counter state-change`	Turns off monitoring state change.

1. The error-pkt-port counter, the err-pkt-from xbar counter, and the err-pkt-to-xbar counter are all ASIC counters. All ASIC counters are turned off by default. The Asic counters are queried every 10 seconds. If the ASIC corresponding to a specific ASIC counter sends or receives any error packets during a 10-second interval, an error frame event occurs during the interval for that counter.
2. Falling-threshold value should be less than the rising-threshold value.

Activating a Port Monitor Policy

Detailed Steps

To activate a port monitor policy, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-monitor activate` *policyname*	Activates the specified port monitor policy.
	`switch(config)# port-monitor activate`	Activates the default port monitor policy.
	`switch(config)# no port-monitor activate` *policyname*	Deactivates the specified port monitoring policy.

Configuring a Port Monitor Port Guard

Detailed Steps

To configure a port monitor policy, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-monitor name` *policyname*	Specifies the policy name and enters the port monitoring policy configuration mode.
	`switch(config)# no port-monitor name` *policyname*	Removes the policy.
Step 3	`switch(config-port-monitor)# counter link-loss poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold` *percentage2* `event` *event-id* `portguard flap`	Specifies the delta link loss, poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered. It also specifies that the port is flappped (is up or down) when the event occurs.
Step 4	`switch(config-port-monitor)# counter link-loss poll-interval` *seconds* `delta rising-threshold` *percentage1* `event` *event-id* `falling-threshold` *percentage2* `event` *event-id* `portguard errordisable`	Specifies the delta link loss, poll interval in seconds, the thresholds in percentage, and the event IDs of events to be triggered. It also specifies that the interface is down (error disabled) when the event occurs.

Note By default, the port monitor port guard is disabled. To enable this feature, you must explicitly configure the port monitor port guard feature on a particular counter by performing Step 3 or Step 4.

Configuring Port Group Monitor

This section includes the following topics:

- Enabling Port Group Monitor, page 2-35
- Configuring a Port Group Monitor Policy, page 2-35

- Reverting to the Default Policy for a Specific Counter, page 2-36
- Turning Off the Monitoring of Specific Counter, page 2-36
- Activating a Port Group Monitor Policy, page 2-37

Enabling Port Group Monitor

Detailed Steps

To enable port group monitor, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-group-monitor enable`	Enables (default) port group monitoring.
	`switch(config)# no port-group-monitor enable`	Disables port group monitoring.

Configuring a Port Group Monitor Policy

Detailed Steps

To configure port group monitor policy, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-group-monitor name` *policyname*	Specifies the policy name and enters the port group monitoring policy configuration mode.
	`switch(config)# no port-group-monitor name` *policyname*	Removes the policy.
Step 3	`switch(config-port-group-monitor)# counter rx-performance poll-interval` *seconds* `delta rising-threshold` *percentage1* `falling-threshold` *percentage2*	Specifies the delta Rx counter poll interval in seconds and thresholds in percentage.
	`switch(config-port-group-monitor)# counter tx-performance poll-interval` *seconds* `delta rising-threshold` *percentage1* `falling-threshold` *percentage2*	Specifies the delta Tx counter poll interval in seconds and thresholds in percentage.
	`switch(config-port-group-monitor)# no counter tx-performance`	[1]Reverts to the [2]default policy.
Step 4	`switch(config-port-group-monitor)# monitor counter rx-performance`	Turns on Rx performance monitoring.
	`switch(config-port-group-monitor)# monitor counter tx-performance`	Turns on Tx performance monitoring.
	`switch(config-port-group-monitor)# no monitor counter tx-performance`	[3]Turns off Tx performance monitoring.

1. See Reverting to the Default Policy for a Specific Counter, page 2-36.
2. See Port Group Monitor, page 2-16
3. See Turning Off the Monitoring of Specific Counter, page 2-36.

Reverting to the Default Policy for a Specific Counter

When the **no counter** command is used in the **config-port-group-monitor** mode, the specified counter polling values will revert to the default values as seen in the following example:

```
switch(config)# port-group-monitor name PGMON_policy
switch(config-port-group-monitor)# counter tx-performance poll-interval 100 delta
rising-threshold 65 falling-threshold 25
switch(config)# show port-group-monitor PGMON_policy

Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ---------------------
RX Performance   Delta      60       80                   10
TX Performance   Delta      100      65                   25

switch(config)# port-group-monitor name PGMON_policy
switch(config-port-group-monitor)# no counter tx-performance
switch(config)# show port-group-monitor PGMON_policy


Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ---------------------
RX Performance   Delta      60       80                   10
TX Performance   Delta      60      80                   10
--------------------------------------------------------------------------------
```

Turning Off the Monitoring of Specific Counter

When the **no monitor counter** command is used in the **config-port-group-monitor** mode, it turns off the monitoring of the specified counter in the given policy as seen in the following example:

```
switch(config)# show port-group-monitor PGMON_policy

Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ---------------------
RX Performance   Delta      26       450                  250
TX Performance   Delta      60       100                  80
--------------------------------------------------------------------------------

switch(config)# port-group-monitor name PGMON_policy
switch(config-port-group-monitor)# no monitor counter rx-performance

switch(config)# show port-group-monitor PGMON_policy

Policy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
```

```
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      26       450                  250
TX Performance   Delta      60       100                  80
--------------------------------------------------------------------------------
```

Activating a Port Group Monitor Policy

To activate a port group monitor policy, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# port-group-monitor activate` *policyname*	Activates the specified port group monitor policy.
	`switch(config)# port-group-monitor activate`	Activates the default port group monitor policy.
	`switch(config)# no port-group-monitor activate` *policyname*	Deactivates the specified port group monitor policy.

Configuring Management Interfaces

Detailed Steps

To configure the mgmt0 Ethernet interface to connect over IPv4, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Selects the management Ethernet interface on the switch and enters interface configuration submode.
Step 3	`switch(config-if)# ip address 10.16.1.2 255.255.255.0`	Configures the IPv4 address and IPv4 subnet mask.
Step 4	`switch(config-if)# no shutdown`	Enables the interface.
Step 5	`switch(config-if)# exit` `switch(config)#`	Returns to configuration mode.
Step 6	`switch(config)# ip default-gateway 1.1.1.4`	Configures the default gateway IPv4 address.
Step 7	`switch(config)# exit` `switch#`	Returns to EXEC mode.
Step 8	`switch# copy running-config startup-config`	(Optional) Saves your configuration changes to the file system. **Note** If you want to save your configuration, you can enter this command at any time.

To configure the mgmt0 Ethernet interface to connect over IPv6, follow these steps:

	Command	Purpose
Step 1	`switch# config terminal` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface mgmt0` `switch(config-if)#`	Selects the management Ethernet interface on the switch and enters interface configuration submode.
Step 3	`switch(config-if)# ipv6 enable`	Enables IPv6 and assigns a link-local address on the interface.
Step 4	`switch(config-if)# ipv6 address ipv6 address 2001:0db8:800:200c::417a/64`	Specifies an IPv6 unicast address and prefix length on the interface.
Step 5	`switch(config-if)# no shutdown`	Enables the interface.
Step 6	`switch(config-if)# end` `switch#`	Returns to EXEC mode.
Step 7	`switch# copy running-config startup-config`	(Optional) Saves your configuration changes to the file system. **Note** If you want to save your configuration, you can enter this command at any time.

Creating VSAN Interfaces

Detailed Steps

To create a VSAN interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface vsan 2` `switch(config-if)#`	Configures a VSAN with the ID 2.
Step 3	`switch(config-if)# no shutdown`	Enables the VSAN interface.

Configuring Slow Drain Device Detection and Congestion Avoidance

This section includes the following topics:

- Configuring Congestion Frame Timeout Value, page 2-39
- Configuring Stuck Frame Timeout Value, page 2-39
- Configuring No-Credit Timeout Value, page 2-39
- Configuring Credit Loss Recovery Threshold and Action, page 2-40
- Configuring Average Credit Nonavailable Duration Threshold and Action, page 2-41

Configuring Congestion Frame Timeout Value

The default congestion frame timeout value is 500 milliseconds. We recommend that you retain the default configuration for the ISLs and configure a value that does not exceed the default value for the edge ports. If the frame is in the switch for a longer time than the configured congestion frame timeout, it gets dropped, which empties the buffer space in the ISL and alleviates the congestion.

To configure the congestion frame timeout value, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system default interface congestion timeout` milliseconds `mode {core | edge}`	Configures a new congestion frame timeout value in milliseconds and the port mode for the device.
Step 3	`switch(config)# system default interface congestion mode {core | edge}`	Configures the default congestion frame timeout value in milliseconds and the port mode for the device. The congestion timeout range is 100-1000 milliseconds.
Step 4	`switch# show logging onboard flow-control request-timeout`	(Optional) Displays the request timeout for a source-destination pair per module with the timestamp information.

Configuring Stuck Frame Timeout Value

The default stuck frame timeout value is 500 ms. We recommend that you retain the default configuration for ISLs and configure a value not exceeding 500 ms (100 to 200 ms) for fabric F ports.

To configure the stuck frame timeout value, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system timeout congestion-drop` *seconds* `mode E|F`	Specifies the stuck frame timeout value in ms and the port mode for the switch.
	`switch(config)# system timeout congestion-drop default mode E|F`	Specifies the default stuck frame timeout port mode for the switch.

Configuring No-Credit Timeout Value

When the port does not have the credits for the configured period, no-credit timeout can be enabled on that port. This will result in all frames coming to that port getting dropped in the egress. This will free the buffer space in the ISL link, which carries traffic for this port. This will help reduce fabric slow down and congestion on other unrelated flows using the same link.

The frames that will be dropped would have just entered the switch or would have stayed in the switch for the configured timeout value. These are preemptive drops and will clear the congestion completely compared to the stuck frame timeout value.

No-credit timeout feature is disabled by default. We recommend that you retain the default configuration for ISLs and configure a value not exceeding 500 ms (200 to 300 ms) for fabric F ports.

Note The no-credit timeout value and stuck frame timeout value are interlinked. The no-credit timeout value must always be greater than the stuck frame timeout value.

To configure the no-credit timeout value, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# system timeout no-credit-drop` *seconds* `mode E|F`	Specifies the no-credit timeout value and port mode for the switch.
	`switch(config)# system timeout no-credit-drop default mode E|F`	Specifies the default no-credit timeout value port mode for the switch.

Configuring Credit Loss Recovery Threshold and Action

When the port detects the credit loss condition and recovers, then the port can be error-disabled, a trap can be sent with interface details, and a syslog can be generated with interface details. When the configured threshold is exceeded, one or more of these actions can be combined together. These actions can be turned on or off depending on situation. The port monitor feature provides the command line interface to configure the thresholds and action.

The thresholds are that the credit loss recovery can be between 1 and 10 and the interval can be 1 second to 1 hour. The default value is 3 in 10 minutes and generates a syslog.

When the port sees the credit loss condition and fails to recover, the port flaps. This function is already part of port guard, and you can configure the supported actions using the Port Guard feature.

To configure credit loss recovery threshold and action, refer to the following example. The following example shows the credit loss recovery threshold and action configuration:

```
switch# show port-monitor
Policy Name  : Cisco
Admin status : Active
Oper status  : Active
Port type    : All Ports
--------------------------------------------------------------------------------
Counter                  Threshold  Interval Rising Threshold event Falling Threshold
event Portguard
-------- ---------  -------- ---------------- ----- ------------------ ----- --------- ---
Link Loss               Delta      60       5               4     1                 4
Not enabled
Sync Loss               Delta      60       5               4     1                 4
Not enabled
Protocol Error          Delta      60       1               4     0                 4
Not enabled
Signal Loss             Delta      60       5               4     1                 4
Not enabled
Invalid Words           Delta      60       1               4     0                 4
Not enabled
Invalid CRC's           Delta      60       5               4     1                 4
Not enabled
RX Performance          Delta      60       2147483648      4     524288000         4
Not enabled
TX Performance          Delta      60       2147483648      4     524288000         4
Not enabled
TX Discards             Delta      60       200             4     10                4
Not enabled
```

```
LR RX                    Delta      60       5               4     1                  4
Not enabled
LR TX                    Delta      60       5               4     1                  4
Not enabled
Timeout Discards         Delta      60       200             4     10                 4
Not enabled
Credit Loss Reco         Delta      60       1               4     0                  4
Not enabled
TX Credit Not Available  Delta      60       10              4     0                  4
Not enabled
--------------------------------------------------------------------------------
```

The following default port monitor policy will be active when the switch comes up:

```
Policy Name  : slowdrain
Admin status : Not Active
Oper status  : Not Active
Port type    : All Ports
--------------------------------------------------------------------------------
----------
Counter                  Threshold  Interval Rising Threshold event Falling Threshold
event Portguard
-------                  ---------  -------- ---------------- ----- ------------------
----- ---------     ------
Credit Loss Reco         Delta      5        4               4     1                  4
Not enabled
TX Credit Not Available  Delta      1        20              4     10                 4
Not enabled
--------------------------------------------------------------------------------
```

Configuring Average Credit Nonavailable Duration Threshold and Action

When the average credit nonavailable duration exceeds the set threshold, the port can be error-disabled, a trap can be sent with interface details, and a syslog can be generated with interface details. One or more of these actions can also be combined together. These actions can be turned on or off depending on the situation. The port monitor feature provides the command line interface to configure the thresholds and action. The threshold configuration can be a percentage of credit nonavailable duration in an interval.

The thresholds are that the credit nonavailable duration can be 0 percent to 100 percent in multiples of 10, and the interval can be 1 second to 1 hour. The default is 10 percent in 1 second and generates a syslog.

To configure average credit nonavailable duration threshold and action, refer to the "Port Monitor" section on page 2-15.

Note This feature is not supported on 1 RU fabric switches.

Verifying Interfaces Configuration

This section includes the following topics:

- Displaying Interface Information, page 2-42
- Displaying TL Port Information, page 2-50
- Displaying the ALPA Cache Contents, page 2-51

- Displaying Port Monitor Status and Policies, page 2-51
- Displaying Port Group Monitor Status and Policies, page 2-53
- Displaying Management Interface Configuration, page 2-55
- Displaying VSAN Interface Information, page 2-55

Displaying Interface Information

The **show interface** command is invoked from the EXEC mode and displays the interface configurations. Without any arguments, this command displays the information for all the configured interfaces in the switch. See Examples 2-3 to 2-10.

Example 2-3 Displays All Interfaces

```
switch# show interface
fc1/1 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:0b:00:05:30:00:8d:de
    Admin port mode is F
    Port mode is F, FCID is 0x610000
    Port vsan is 2
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      134 frames input, 8468 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      154 frames output, 46072 bytes
        0 discards, 0 errors
      1 input OLS, 1 LRR, 0 NOS, 0 loop inits
      1 output OLS, 0 LRR, 1 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
fc1/9 is trunking
    Hardware is Fibre Channel, SFP is long wave laser cost reduced
    Port WWN is 20:09:00:05:30:00:97:9e
    Peer port WWN is 20:0b:00:0b:5f:a3:cc:00
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 255
    Receive B2B Credit is 255
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100,3000)
    Trunk vsans (up)                       (1,100,3000)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 280 bits/sec, 35 bytes/sec, 0 frames/sec
    5 minutes output rate 176 bits/sec, 22 bytes/sec, 0 frames/sec
      4609939 frames input, 8149405708 bytes
        0 discards, 0 errors
```

```
        0 CRC, 0 unknown class
        0 too long, 0 too short
      4638491 frames output, 7264731728 bytes
        0 discards, 0 errors
      3 input OLS, 9 LRR, 1 NOS, 0 loop inits
      9 output OLS, 7 LRR, 1 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
fc1/13 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:0d:00:05:30:00:97:9e
    Admin port mode is auto, trunk mode is on
    Port mode is F, FCID is 0x650100
    Port vsan is 100
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      8696 frames input, 3227212 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      16799 frames output, 6782444 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
sup-fc0 is up
    Hardware is Fibre Channel
    Speed is 1 Gbps
    139597 packets input, 13852970 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    139516 packets output, 16759004 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

You can also specify arguments (a range of interfaces or multiple, specified interfaces) to display interface information. You can specify a range of interfaces by issuing a command with the following example format:

```
interface fc1/1 - 5 , fc2/5 - 7
```

Note The spaces are required before and after the dash (**-**) and before and after the comma (**,**).

Example 2-4 Displays Multiple, Specified Interfaces

```
switch# show interface fc3/13 , fc3/16
fc3/13 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:8d:00:05:30:00:97:9e
    Admin port mode is FX
    Port mode is F, FCID is 0x7b0300
    Port vsan is 1
    Speed is 2 Gbps
    Transmit B2B Credit is 3
```

```
    Receive B2B Credit is 12
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1856 frames input, 116632 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      1886 frames output, 887712 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 1 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.

fc3/16 is up
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:90:00:05:30:00:97:9e
    Admin port mode is FX
    Port mode is F, FCID is 0x7d0100
    Port vsan is 3000
    Speed is 2 Gbps
    Transmit B2B Credit is 3
    Receive B2B Credit is 12
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 504 bits/sec, 63 bytes/sec, 0 frames/sec
    5 minutes output rate 520 bits/sec, 65 bytes/sec, 0 frames/sec
      47050 frames input, 10311824 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      62659 frames output, 10676988 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      1 output OLS, 1 LRR, 0 NOS, 1 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

Example 2-5 Displays a Specific Interface

```
switch# show interface fc2/2
fc2/2 is trunking
    Port description is Trunk to Core-4
    Hardware is Fibre Channel, SFP is short wave laser
    Port WWN is 20:42:00:05:30:00:97:9e
    Peer port WWN is 20:cc:00:05:30:00:50:9e
    Admin port mode is E, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Transmit B2B Credit is 255
    Receive B2B Credit is 255
    Receive data field Size is 2112
    Beacon is turned off
    Belongs to port-channel 2
    Trunk vsans (admin allowed and active) (1,100,3000)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 (100,3000)
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 32 bits/sec, 4 bytes/sec, 0 frames/sec
```

```
      2214834 frames input, 98673588 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      2262415 frames output, 343158368 bytes
        0 discards, 0 errors
      1 input OLS, 1 LRR, 1 NOS, 0 loop inits
      2 output OLS, 1 LRR, 0 NOS, 0 loop inits
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
```

Example 2-6 Displays Port Description

```
switch# show interface description
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
fc3/1              test intest
fc3/2              --
fc3/3              --
fc3/4              TE port
fc3/5              --
fc3/6              --
fc3/10             Next hop switch 5
fc3/11             --
fc3/12             --
fc3/16             --
--------------------------------------------------------------------------------
Interface          Description
--------------------------------------------------------------------------------
port-channel 1     --
port-channel 5     --
port-channel 6     --
```

Example 2-7 Displays Interface Information in a Brief Format

```
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP   Oper  Oper   Port
                 Mode   Trunk                         Mode  Speed  Channel
                        Mode                                (Gbps)
--------------------------------------------------------------------------------
fc1/1      1      E      on      trunking         swl    TE      2    1
fc1/2      1      E      on      trunking         swl    TE      2    1
fc1/3      1      auto   on      SFPAbsent       --     --           --
fc1/4      1      auto   on      SFPAbsent       --     --           --
fc1/5      3000   auto   on      up               swl    F       2    --
...
fc2/2      1      E      on      trunking         swl    TE      2    2
fc2/3      1      auto   on      down             c1610  --           --
fc2/4      1      auto   on      down             c1590  --           --
fc2/5      3000   auto   on      notConnected     lwcr   --           --
fc2/6      1      auto   on      SFPAbsent       --     --           --
...
fc3/16     3000   FX     --      up               swl    F       2    --
fc3/17     1      FX     --      SFPAbsent       --     --           --
...
--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed        MTU
--------------------------------------------------------------------------------
GigabitEthernet4/1      SFPAbsent --                auto         1500
```

```
...
GigabitEthernet4/6      down       10.1.1.2/8        auto         3000
GigabitEthernet4/7      down       10.1.1.27/24      auto         1500
GigabitEthernet4/8      down       --                auto         1500

--------------------------------------------------------------------------------
Interface           Status              Oper Mode           Oper Speed
                                                            (Gbps)
--------------------------------------------------------------------------------
iscsi4/1            down                --
...
--------------------------------------------------------------------------------
Interface          Status                            Speed
                                                     (Gbps)
--------------------------------------------------------------------------------
sup-fc0            up                                1

--------------------------------------------------------------------------------
Interface               Status     IP Address        Speed        MTU
--------------------------------------------------------------------------------
mgmt0                   up         172.19.48.96/25   100 Mbps     1500

--------------------------------------------------------------------------------
Interface               Vsan     Admin    Status            Oper   Oper
                                 Trunk                      Mode   Speed
                                 Mode                              (Gbps)
--------------------------------------------------------------------------------
port-channel 1          1        on       trunking          TE     4
port-channel 2          1        on       trunking          TE     4

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status       Oper  Profile  Port-channel
                  Mode   Trunk               Mode
                         Mode
--------------------------------------------------------------------------------
fcip10     1      auto   on      notConnected --     10       --
```

Example 2-8 Displays Interface Counters

```
switch# show interface counters
fc3/1
    5 minutes input rate 24 bits/sec, 3 bytes/sec, 0 frames/sec
    5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
    3502 frames input, 268400 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    3505 frames output, 198888 bytes
      0 discards
    1 input OLS, 1 LRR, 1 NOS, 0 loop inits
    2 output OLS, 1 LRR, 1 NOS, 0 loop inits
    1 link failures, 1 sync losses, 1 signal losses
.
.
.
fc9/8
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
```

```
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
      16 receive B2B credit remaining
      3 transmit B2B credit remaining.
. . .
sup-fc0
    114000 packets input, 11585632 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    113997 packets output, 10969672 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors

mgmt0
    31557 packets input, 2230860 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    26618 packets output, 16824342 bytes, 0 underruns
      0 output errors, 0 collisions, 7 fifo
      0 carrier errors

vsan1
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
.
.
.
port-channel 1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    0 frames input, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards, 0 CRC, 0 unknown class
      0 too long, 0 too short
    0 frames output, 0 bytes
      0 class-2 frames, 0 bytes
      0 class-3 frames, 0 bytes
      0 class-f frames, 0 bytes
      0 discards
    0 input OLS, 0 LRR, 0 NOS, 0 loop inits
    0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    0 link failures, 0 sync losses, 0 signal losses
```

Note Interfaces 9/8 and 9/9 are not trunking ports and display class 2, 3, and F information as well.

Example 2-9 Displays Interface Counters in Brief Format

```
switch# show interface counters brief

-------------------------------------------------------------------------------
Interface          Input (rate is 5 min avg)      Output (rate is 5 min avg)
                   ----------------------------   ----------------------------
                   Rate     Total                 Rate     Total
```

```
                  Mbits/s  Frames                Mbits/s  Frames
--------------------------------------------------------------------------------
fc3/1             0        3871                  0        3874
fc3/2             0        3902                  0        4232
fc3/3             0        3901                  0        4138
fc3/4             0        3895                  0        3894
fc3/5             0        3890                  0        3897
fc9/8             0        0                     0        0
fc9/9             0        5                     0        4
fc9/10            0        4186                  0        4182
fc9/11            0        4331                  0        4315

--------------------------------------------------------------------------------
Interface         Input (rate is 5 min avg)      Output (rate is 5 min avg)
                  ---------------------------  ---------------------------
                  Rate     Total                Rate     Total
                  Mbits/s  Frames               Mbits/s  Frames
--------------------------------------------------------------------------------
port-channel 1    0        0                    0        0
port-channel 2    0        3946                 0        3946
```

Note The **show interface transceiver** command can only be issued on a switch in the Cisco MDS 9100 Series if the SFP is present (see Example 2-10).

Example 2-10 Displays Transceiver Information

```
switch# show interface transceiver
fc1/1 SFP is present
    name is CISCO-AGILENT
    part number is QFBR-5796L
    revision is
    serial number is A00162193
    fc-transmitter type is short wave laser
    cisco extended id is unknown (0x0)
...
fc1/9 SFP is present
    name is FINISAR CORP.
    part number is FTRJ-1319-7D-CSC
    revision is
    serial number is H11A6ER
    fc-transmitter type is long wave laser cost reduced
    cisco extended id is unknown (0x0)
...
```

Example 2-11 displays the entire running configuration with information for all interfaces. The interfaces have multiple entries in the configuration files to ensure that the interface configuration commands execute in the correct order when the switch reloads.

Example 2-11 Displays the Running Configuration for All Interfaces

```
switch# show running-config
...
interface fc9/1
  switchport speed 2000
...
interface fc9/1
  switchport mode E
...
interface fc9/1
```

```
  channel-group 11 force
  no shutdown
```

Example 2-12 displays the running configuration information for a specified interface. The interface configuration commands are grouped together.

Example 2-12 Displays the Running Configuration for a Specified Interface

```
switch# show running-config interface fc1/1
interface fc9/1
  switchport speed 2000
  switchport mode E
  channel-group 11 force
  no shutdown
```

Example 2-13 displays the running configuration after the **system default switchport mode F** command is executed.

Example 2-13 Displays the Running Configuration After the System Default Switchport Mode F Command is Executed

```
switch# show running-config
version 3.1(3)
system default switchport mode F
interface fc4/1
interface fc4/2
interface fc4/3
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
interface fc4/8
interface fc4/9
interface fc4/10
```

Example 2-14 displays the running configuration after two interfaces are individually configured for mode FL.

Example 2-14 Displays the Running Configuration After Two Interfaces Are Individually Configured for Mode FL

```
switch# show running-config
version 3.1(3)
system default switchport mode F
interface fc4/1
  switchport mode FL
interface fc4/2
interface fc4/3
  switchport mode FL
interface fc4/4
interface fc4/5
interface fc4/6
interface fc4/7
interface fc4/8
interface fc4/9
interface fc4/1
```

Example 2-15 displays interface information in a brief format after the **system default switchport mode F** command is executed. Example 2-16 displays interface information in a brief format after two interfaces are individually configured for mode FL.

Example 2-15 Displays Interface Information in a Brief Format After the System Default Switchport Mode F Command is Executed

```
switch# show interface brief
--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                          (Gbps)
--------------------------------------------------------------------------------
fc4/1      1      F      --      notConnected     swl    --           --
fc4/2      1      F      --      notConnected     swl    --           --
fc4/3      1      F      --      notConnected     swl    --           --
fc4/4      1      F      --      notConnected     swl    --           --
fc4/5      1      F      --      sfpAbsent        --     --           --
fc4/6      1      F      --      sfpAbsent        --     --           --
fc4/7      1      F      --      sfpAbsent        --     --           --
fc4/8      1      F      --      sfpAbsent        --     --           --
fc4/9      1      F      --      sfpAbsent        --     --           --
```

Example 2-16 Displays Interface Information in a Brief Format After Two Interfaces Are Individually Configured for Mode FL

```
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                          (Gbps)
--------------------------------------------------------------------------------
fc4/1      1      FL     --      notConnected     swl    --           --
fc4/2      1      F      --      notConnected     swl    --           --
fc4/3      1      FL     --      notConnected     swl    --           --
fc4/4      1      F      --      notConnected     swl    --           --
fc4/5      1      F      --      sfpAbsent        --     --           --
fc4/6      1      F      --      sfpAbsent        --     --           --
fc4/7      1      F      --      sfpAbsent        --     --           --
fc4/8      1      F      --      sfpAbsent        --     --           --
fc4/9      1      F      --      sfpAbsent        --     --           --
fc4/10     1      F      --      sfpAbsent        --     --           --
```

Displaying TL Port Information

The **show tlport** command displays the TL port interface configurations. This command provides a list of all TL ports configured in a switch and shows the associated VSAN, the FC ID for the port (only domain and area are valid), and the current operational state of the TL port (up or initializing). See Example 2-17 through Example 2-20.

Example 2-17 Displays the TL Ports in All VSANs

```
switch# show tlport list
-------------------------------
Interface Vsan FC-ID    State
------------------------ ------
fc1/16    1    0x420000 Init
fc2/26    1    0x150000 Up
```

TL ports allow a private device (devices that physically reside on the loop) to see a fabric device and vice-versa by proxying fabric devices on the loop. Fabric devices are proxied by allocating each fabric device an ALPA on this loop.

In addition to these proxied devices, other virtual devices (local or remote domain controller addresses) are also allocated ALPAs on the loop. A switch reserves the ALPA for its own communication with private devices, and the switch acts as a SCSI initiator.

The first column in the output of the **show tlport interface** command is the ALPA identity of the device on the loop. The columns that follow include the port WWNs, the node WWNs for each device, the device as a SCSI initiator or target, and the real FC ID of the device.

Example 2-18 Displays the Detailed Information for a Specific TL Port

```
switch# show tlport interface fc1/16 all
fc1/16 is up, vsan 1, FCID 0x420000
--------------------------------------------------------------------------------
alpa pWWN                    nWWN                    SCSI Type Device  FC-ID
--------------------------------------------------------------------------------
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Proxied 0xfffc42
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target    Private 0x420073
0xef 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator Switch  0x0000ef
```

Example 2-19 Displays TL Port Information for Private Devices

```
switch# show tlport interface fc 1/16 private
fc1/16 is up, vsan 1, FCID 0x420000
------------------------------------------------------------------------
alpa pWWN                    nWWN                    SCSI Type FC-ID
------------------------------------------------------------------------
0x73 22:00:00:20:37:39:ae:54 20:00:00:20:37:39:ae:54 Target    0x420073
0x74 22:00:00:20:37:38:d3:de 20:00:00:20:37:38:d3:de Target    0x420074
```

Example 2-20 Displays TL Port Information for Proxied Devices

```
switch# show tlport interface fc 1/16 proxied
fc1/16 is up, vsan 1, FCID 0x420000
------------------------------------------------------------------------
alpa pWWN                    nWWN                    SCSI Type FC-ID
------------------------------------------------------------------------
0x01 20:10:00:05:30:00:4a:de 20:00:00:05:30:00:4a:de Initiator 0xfffc42
0x02 21:00:00:e0:8b:01:95:e7 20:00:00:e0:8b:01:95:e7 Initiator 0x420100
```

Displaying the ALPA Cache Contents

The **show tlport alpa-cache** command displays the contents of the ALPA cache.

```
switch# show tlport alpa-cache
-----------------------------------------------------------
alpa              pWWN               Interface
-----------------------------------------------------------
0x02  22:00:00:20:37:46:09:bd     fc1/2
0x04  23:00:00:20:37:46:09:bd     fc1/2
```

The first entry indicates that if a device with a pWWN of 22:00:00:20:37:46:09:bd is exported on TL port fc1/2, then the pWWN is allocated an alpa 0x02 (if available).

Displaying Port Monitor Status and Policies

The following commands display information regarding port monitor:

```
switch# show port-monitor status
Port Monitor    : Enabled
Active Policies : sample
Last 10 logs :

switch# show port-monitor
--------------------------------------------------------------------------------
Port Monitor : enabled
--------------------------------------------------------------------------------
Policy Name  : sample
Admin status : Not Active
Oper status  : Not Active
Port type    : All Ports
--------------------------------------------------------------------------------
Counter          Threshold  Interval Rising Threshold event Falling Threshold  event
portguard
-------          ---------  -------- ---------------- ----- ------------------ ----- -----
Link Loss        Delta      60       5                4     1                  4
Not enabled
Sync Loss        Delta      60       5                4     1                  4
Not enabled
Protocol Error   Delta      60       1                4     0                  4
Not enabled
Signal Loss      Delta      60       5                4     1                  4
Not enabled
Invalid Words    Delta      60       1                4     0                  4
Not enabled
Invalid CRC's    Delta      60       5                4     1                  4
Not enabled
RX Performance   Delta      60       2147483648       4     524288000          4
Not enabled
TX Performance   Delta      60       2147483648       4     524288000          4
Not enabled
--------------------------------------------------------------------------------

Policy Name  : default
Admin status : Not Active
Oper status  : Not Active
Port type    : All Ports
--------------------------------------------------------------------------------
Counter          Threshold  Interval Rising Threshold event Falling Threshold  event
portguard
-------          ---------  -------- ---------------- ----- ------------------ ----- -----
Link Loss        Delta      60       5                4     1                  4
Not enabled
Sync Loss        Delta      60       5                4     1                  4
Not enabled
Protocol Error   Delta      60       1                4     0                  4
Not enabled
Signal Loss      Delta      60       5                4     1                  4
Not enabled
Invalid Words    Delta      60       1                4     0                  4
Not enabled
Invalid CRC's    Delta      60       5                4     1                  4
Not enabled
RX Performance   Delta      60       2147483648       4     524288000          4
Not enabled
TX Performance   Delta      60       2147483648       4     524288000          4
Not enabled
--------------------------------------------------------------------------------

switch# show port-monitor active
Policy Name  : sample
Admin status : Active
```

```
Oper status  : Active
Port type    : All Access Ports
--------------------------------------------------------------------------------
Counter          Threshold  Interval Rising Threshold event Falling Threshold  event
portguard
-------          ---------  -------- ---------------- ----- ------------------ ----- -----
Link Loss        Delta      60       5                4     1                  4
enabled
Sync Loss        Delta      60       5                4     1                  4
enabled
Protocol Error   Delta      60       1                4     0                  4
enabled
Signal Loss      Delta      60       5                4     1                  4
enabled
Invalid Words    Delta      60       1                4     0                  4
enabled
Invalid CRC's    Delta      60       5                4     1                  4
enabled
RX Performance   Delta      60       2147483648       4     524288000          4
enabled
TX Performance   Delta      60       2147483648       4     524288000          4
enabled
--------------------------------------------------------------------------------

switch# show port-monitor sample
Policy Name  : sample
Admin status : Active
Oper status  : Active
Port type    : All Access Ports
--------------------------------------------------------------------------------
Counter          Threshold  Interval Rising Threshold event Falling Threshold  event
portgurard
-------          ---------  -------- ---------------- ----- ------------------ ----- -----
Link Loss        Delta      60       5                4     1                  4
enabled
Sync Loss        Delta      60       5                4     1                  4
enabled
Protocol Error   Delta      60       1                4     0                  4
enabled
Signal Loss      Delta      60       5                4     1                  4
enabled
Invalid Words    Delta      60       1                4     0                  4
enabled
Invalid CRC's    Delta      60       5                4     1                  4
enabled
RX Performance   Delta      60       2147483648       4     524288000          4
enabled
TX Performance   Delta      60       2147483648       4     524288000          4
enabled
--------------------------------------------------------------------------------
```

Displaying Port Group Monitor Status and Policies

The following commands display information about port group monitor:

```
switch# show port-group-monitor status
Port Group Monitor : Enabled
Active Policies : pgm2
Last 10 logs :
switch#
```

```
switch# show port-group-monitor
--------------------------------------------------------------------------------
Port Group Monitor : enabled
--------------------------------------------------------------------------------
Policy Name  : pgm1
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      60       50                   10
TX Performance   Delta      60       50                   10
--------------------------------------------------------------------------------
Policy Name  : pgm2
Admin status : Active
Oper status  : Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      60       80                   10
TX Performance   Delta      60       80                   10
--------------------------------------------------------------------------------
Policy Name  : default
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      60       80                   20
TX Performance   Delta      60       80                   20
--------------------------------------------------------------------------------

switch# show port-group-monitor active
Policy Name  : pgm2
Admin status : Active
Oper status  : Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      60       80                   10
TX Performance   Delta      60       80                   10
--------------------------------------------------------------------------------

switch# show port-group-monitor PGMON_policy
PPolicy Name  : PGMON_policy
Admin status : Not Active
Oper status  : Not Active
Port type    : All Port Groups
--------------------------------------------------------------------------------
Counter          Threshold  Interval %ge Rising Threshold %ge Falling Threshold
-------          ---------  -------- -------------------- ----------------------
RX Performance   Delta      26       450                  250
TX Performance   Delta      60       100                  80
--------------------------------------------------------------------------------
```

Displaying Management Interface Configuration

To display the management interface configuration, use the **show interface mgmt 0** command.

```
switch# show interface mgmt 0
mgmt0 is up
    Hardware is FastEthernet
    Address is 000c.30d9.fdbc
    Internet address is 10.16.1.2/24
    MTU 1500 bytes, BW 100 Mbps full Duplex
    26388 packets input, 6101647 bytes
      0 multicast frames, 0 compressed
      0 input errors, 0 frame, 0 overrun 0 fifo
    10247 packets output, 2389196 bytes, 0 underruns
      0 output errors, 0 collisions, 0 fifo
      0 carrier errors
```

Displaying VSAN Interface Information

To display VSAN interface information, use the **show interface vsan** command.

```
switch# show interface vsan 2
vsan2 is up, line protocol is up
    WWPN is 10:00:00:05:30:00:59:1f, FCID is 0xb90100
    Internet address is 10.1.1.1/24
    MTU 1500 bytes, BW 1000000 Kbit
    0 packets input, 0 bytes, 0 errors, 0 multicast
    0 packets output, 0 bytes, 0 errors, 0 dropped
```

Configuring Fibre Channel Interfaces

- Information About Fibre Channel Interfaces, page 3-1
- Guidelines and Limitations, page 3-18
- Default Settings, page 3-25
- Configuring Fibre Channel Interfaces, page 3-26
- Verifying Fibre Channel Interfaces Configuration, page 3-43
- Configuration Examples for Fibre Channel Interfaces, page 3-45

Information About Fibre Channel Interfaces

This section includes the following topics:

- Generations of Modules and Switches, page 3-1
- Port Groups, page 3-3
- Port Rate Modes, page 3-5
- Port Speed, page 3-10
- Dynamic Bandwidth Management, page 3-10
- Out-of-Service Interfaces, page 3-11
- Oversubscription Ratio Restrictions, page 3-11
- Bandwidth Fairness, page 3-17

Generations of Modules and Switches

Cisco MDS 9000 Family hardware modules and switches are categorized into generations based on the time of introduction, capabilities, features, and compatibilities:

- Generation 1—Modules and switches with a maximum port speed of 2 Gbps.
- Generation 2—Modules and switches with a maximum port speed of 4 Gbps.
- Generation 3—Modules and switches with a maximum port speed of 8 Gbps.
- Generation 4—Modules with a maximum port speed of 8-Gbps or 10-Gbps.

The Cisco MDS 9500 Series switches, Cisco MDS 9222i, Cisco MDS 9216A, and Cisco MDS 9216i switches support the Generation 2 modules. Each module or switch can have one or more ports in port groups that share common resources such as bandwidth and buffer credits.

In addition to supporting Generation 2 modules, the Cisco MDS 9500 Series switches and the Cisco MDS 9222i switch support the Generation 3 modules. Similar to Generation 2, each Generation 3 or Generation 4 module can have one or more ports in port groups that share common resources such as bandwidth and buffer credits.

Generation 3 modules are supported on the Cisco MDS 9506 and 9509 switches with Supervisor-2 modules. The MDS 9513 Director supports 4/44-port Host-Optimized Fibre Channel switching module with either Fabric 1 or Fabric 2 modules, but requires Fabric 2 module for support of the 48-port and the 24-port 8-Gbps Fibre Channel switching modules. The MDS 9222i switch supports the 4/44-port Host-Optimized Fibre Channel switching module.

The Cisco 9500 Series switches support the following Generation 4 modules: the 48-port 8-Gbps Advanced Fibre Channel switching module (DS-X9248-256K9) and the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9). Cisco MDS NX-OS Release 5.2(1) or higher is required to support the Generation 4 modules.

Table 3-1 identifies the Generation 2, Generation 3, and Generation 4 modules, as well as the Fabric switches.

Table 3-1 Fibre Channel Modules and Fabric Switches

Part Number	Product Name and Description
Generation 4 Modules	
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module.
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module.
DS-X9530-SF2A-K9	Supervisor-2A module for Cisco MDS 9500 Series switches.
DS-13SLT-FAB3	Fabric 3 module that enables the 32-port and the 48-port 8-Gbps Advanced Fibre Channel switching module to use the full 96-Gbps or 256-Gbps backplane crossbar bandwidth.
Generation 3 Modules	
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module.
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module.
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module
DS-13SLT-FAB2	Fabric 2 module that enables the 24-port and the 48-port 8-Gbps Fibre Channel switching module to use the full 96-Gbps backplane bandwidth with any-to-any connectivity.
Generation 3 Fabric Switch	
DS-C9148-K9	Cisco MDS 9148 Fabric switch. 48-port 8-Gbps Fabric switch.
Generation 2 Modules	
DS-X9148	48-port 4-Gbps Fibre Channel switching module.
DS-X9124	24-port 4-Gbps Fibre Channel switching module.
DS-X9304-18K9	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports.

Table 3-1 Fibre Channel Modules and Fabric Switches (continued)

Part Number	Product Name and Description
DS-X9112	12-port 4-Gbps Fibre Channel switching module.
DS-X9704	4-port 10-Gbps Fibre Channel switching module.
DS-X9530-SF2-K9	Supervisor-2 module for Cisco MDS 9500 Series switches.
Generation 2 Fabric Switches	
DS-C9134-K9	Cisco MDS 9134 Fabric switch. 32-port 4-Gbps Fabric switch with 2 additional 10-Gbps ports.
DS-C9124-K9	Cisco MDS 9124 Fabric switch. 24-port 4-Gbps Fabric switch.
DS-C9222i-K9	Cisco MDS 9222i Multiservice Modular switch. 18-port 4-Gbps switch with 4-Gigabit Ethernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family switching and services modules.

Note Generation 2 Fibre Channel switching modules are not supported on the Cisco MDS 9216 switch; however, they are supported by both the Supervisor-1 module and the Supervisor-2 module.

For detailed information about the installation and specifications for these modules and switches, refer to the hardware installation guide for your switch.

Port Groups

Each module or switch can have one or more ports in port groups that share common resources such as bandwidth and buffer credits. Port groups are defined by the hardware consisting of sequential ports. For example, ports 1 through 6, ports 7 through 12, ports 13 through 18, ports 19 through 24, ports 25 through 30, 31 through 36, and ports 37 through 42, 43 through 48 are the port groups on the 48-port 8-Gbps Advanced Fibre Channel switching modules.

Table 3-2 shows the bandwidth and number of ports per port group for the Generation 2, Generation 3, and Generation 4 Fibre Channel modules, and Generation 2 and Generation 3 Fabric switches.

Table 3-2 Bandwidth and Port Groups for the Fibre Channel Modules and Fabric Switches

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
Generation 4 Modules				
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module.	6	32.4[1] or 12.8[2]	8 or 10 Gbps—depending on the configuration

Table 3-2 Bandwidth and Port Groups for the Fibre Channel Modules and Fabric Switches (continued)

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module.	4	32.4[1] or 12.8[2]	8 or 10 Gbps—depending on the configuration
Generation 3 Modules				
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module	6	12.8	8 Gbps
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module	3	12.8	8 Gbps
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module	12	12.8	8/4 Gbps[3]
Generation 3 Fabric Switches				
DS-C9148-K9 (Cisco MDS 9148 Fabric switch)	48-port 8-Gbps Fabric switch	4	32	8 Gbps
Generation 2 Modules				
DS-X9148	48-port 4-Gbps Fibre Channel switching module	12	12.8	4 Gbps
DS-X9124	24-port 4-Gbps Fibre Channel switching module	6	12.8	4 Gbps
DS-X9304-18K9 (MSM-18/4 Multiservice module)	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports	6	12.8	4 Gbps
DS-X9112	12-port 4-Gbps Fibre Channel switching module	3	12.8	4 Gbps
DS-X9704	4-port 10-Gbps Fibre Channel switching module	1	10	10 Gbps
Generation 2 Fabric Switches				

Table 3-2 Bandwidth and Port Groups for the Fibre Channel Modules and Fabric Switches (continued)

Part Number	Product Name/ Description	Number of Ports Per Port Group	Bandwidth Per Port Group (Gbps)	Maximum Bandwidth Per Port (Gbps)
DS-C9134-K9 (Cisco MDS 9134 Fabric switch)	32-port 4-Gbps Fabric switch	4	16	4 Gbps
	2-port 10-Gbps Fabric switch	1	10	10 Gbps
DS-C9124K9 (Cisco MDS 9124 Fabric switch)	24-port 4-Gbps Fabric switch	4	16	4 Gbps
DS-C9222i-K9 (Cisco MDS 9222i Multiservice Modular switch)	18-port 4-Gbps, 4 Gigabit Ethernet ports and a modular expansion slot.	6	12.8	4 Gbps

1. This bandwidth is available with the Fabric 3 module (DS-13SLT-FAB3) in the MDS 9513 switch.
2. This bandwidth is available with the Fabric 2 module (DS-13SLT-FAB2) in the MDS 9513 switch, and with the Supervisor-2 (DS-X9530-SF2-K9) or Supervisor-2A module (DS-X9530-SF2AK9) in the MDS 9509 switch or MDS 9506 switch.
3. A maximum of four ports (one per port group) in a 4/44-port 8-Gbps switching module can operate at 8-Gbps bandwidth in dedicated or shared mode. All the other ports can operate at a maximum of 4-Gbps in shared mode or dedicated mode.

Port Rate Modes

In Generation 2, Generation 3, and Generation 4 modules, you can configure the port rate modes. The *port rate mode* configuration is used to determine the bandwidth allocation for ports in a port group. Two port rate modes are supported:

- Dedicated Rate Mode—A port is allocated required fabric bandwidth to sustain line traffic at the maximum operating speed configured on the port. For more information, see the "Dedicated Rate Mode" section on page 3-7.
- Shared Rate Mode—Multiple ports in a port group share data paths to the switch fabric and share bandwidth. For more information, see the "Shared Rate Mode" section on page 3-8.

Note In Generation 1 modules, you cannot configure the port rate modes. The mode is determined implicitly based on the port mode and line card type.

Note Port rate modes are not supported on the Cisco Fabric Switch for HP c-Class Blade System, and the Cisco Fabric Switch for IBM Blade Center.

Table 3-3 shows the modules that support dedicated, shared, and the default rate modes.

Table 3-3 Port Rate Mode Support on Generation 2, Generation 3, and Generation 4 Modules and Switches

Part Number	Product Name/ Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode	Default Speed Mode and Rate Mode on All Ports
Generation 4 Modules				
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module	Yes	Yes[1]	Auto, Shared
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module	Yes	Yes[1]	Auto, Shared
Generation 3 Modules				
DS-X9248-96K9	48-Port 8-Gbps Fibre Channel switching module	Yes	Yes[2]	Auto, Shared
DS-X9224-96K9	24-Port 8-Gbps Fibre Channel switching module	Yes	Yes[2]	Auto, Shared
DS-X9248-48K9	4/44-Port 8-Gbps Host-Optimized Fibre Channel switching module	Yes	Yes[2]	Auto Max 4 Gbps, Shared
Generation 3 Fabric Switches				
DS-C9148-K9 (Cisco MDS 9148 Fabric switch)	48-port 8-Gbps Fabric switch	Yes	No	Auto, Dedicated
Generation 2 Modules				
DS-X9148	48-port 4-Gbps Fibre Channel switching module [3]	Yes	Yes	Auto, Shared
DS-X9124	24-port 4-Gbps Fibre Channel switching module	Yes	Yes	Auto, Shared
DS-X9304-18K9 (MSM-18/4 Multiservice module)	18-port 4-Gbps Fibre Channel switching module with 4-Gigabit Ethernet ports	Yes	Yes	Auto, Shared
DS-X9112	12-port 4-Gbps Fibre Channel switching module	Yes	No	Auto, Dedicated
DS-X9704	4-port 10-Gbps Fibre Channel switching module	Yes	No	Auto, Dedicated
Generation 2 Fabric Switches				

Table 3-3 Port Rate Mode Support on Generation 2, Generation 3, and Generation 4 Modules and Switches (continued)

Part Number	Product Name/ Description	Supports Dedicated Rate Mode	Supports Shared Rate Mode	Default Speed Mode and Rate Mode on All Ports
DS-C9134-K9 (Cisco MDS 9134 Fabric switch)	32-port 4-Gbps Fabric switch	Yes	Yes	Auto, Shared
	2-port 10-Gbps Fabric switch	Yes	No	Auto, Dedicated
DS-C9124-K9 (Cisco MDS 9124 Fabric switch)	24-port 4-Gbps Fabric switch[4]	Yes	No	Auto, Dedicated
DS-C9222i-K9 (Cisco MDS 9222i Multiservice Modular switch)	18-port 4-Gbps Fibre Channel switch with 4-Gigabit Ethernet IP storage services ports, and a modular expansion slot to host Cisco MDS 9000 Family Switching and Services Modules	Yes	Yes	Auto, Shared

1. Supports shared rate mode.
2. Shared rate mode is supported on Fx ports only and no ISLs.
3. All ports in a 48-port 4-Gbps switching module can operate in dedicated rate mode with a 1-Gbps operating speed. However, if you configure one or more ports to operate in 2-Gbps or 4-Gbps dedicated rate mode, some of the other ports in the port group would have to operate in shared mode.
4. All ports in a 24-port 4-Gbps switching module can operate in dedicated rate mode with a 2-Gbps operating speed. However, if you configure one or more ports to operate in 4-Gbps dedicated rate mode, some of the other ports in the port group would have to operate in shared mode.

Dedicated Rate Mode

When port rate mode is configured as dedicated, a port is allocated required fabric bandwidth and related resources to sustain line rate traffic at the maximum operating speed configured for the port. In this mode, ports do not use local buffering and all receive buffers are allocated from a global buffer pool (see the "Buffer Pools" section on page 4-2).

Table 3-4 shows the bandwidth provided by the various port speed configurations on the 8-Gbps Advanced Fibre Channel switching modules.

Table 3-4 Bandwidth Reserved for the Port Speeds on Generation 4 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	8 Gbps
8-Gbps	
Auto with 4-Gbps maximum	4 Gbps
4-Gbps	

Table 3-4 Bandwidth Reserved for the Port Speeds on Generation 4 Switching Modules

Configured Speed	Reserved Bandwidth
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

Table 3-5 shows the bandwidth provided by the various port speed configurations on the 8-Gbps Fibre Channel switching modules.

Table 3-5 Bandwidth Reserved for the Port Speeds on Generation 3 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	8 Gbps
8-Gbps	
Auto with 4-Gbps maximum	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

Table 3-6 shows the amount of bandwidth reserved for a configured port speed on 4-Gbps switching modules.

Table 3-6 Bandwidth Reserved for the Port Speeds on Generation 2 Switching Modules

Configured Speed	Reserved Bandwidth
Auto	4 Gbps
4-Gbps	
Auto with 2-Gbps maximum	2 Gbps
2-Gbps	
1-Gbps	1 Gbps

Note The 4-Port 10-Gbps Fibre Channel module ports in auto mode only support auto speed mode at 10 Gbps.

Shared Rate Mode

When port rate mode is configured as shared, multiple ports within a port group share data paths to the switch fabric so that fabric bandwidth and related resources are shared. Often, the available bandwidth to the switch fabric may be less than the negotiated operating speed of a port. Ports in this mode use local buffering for the BB_credit buffers.

All ports in 8-Gbps Advanced Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8 Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated rate mode with 1-Gbps, 2-Gbps, 4-Gbps, or 8 Gbps operating speed.

All ports in 4-Gbps Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, or 4-Gbps traffic. However, it is possible to configure one or more ports in a port group to operate in dedicated rate mode with 1-Gbps, 2-Gbps, or 4-Gbps operating speed.

All ports in the 32-Port or 48-Port 8-Gbps Advanced Fibre Channel modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8-Gbps traffic in a maximum or 32 or 48 ports.

All ports in the 48-Port and 24-Port 8-Gbps Fibre Channel switching modules where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps, or 8-Gbps traffic.

In the 4/44-Port 8-Gbps Host-Optimized Fibre Channel switching module, all the ports where bandwidth is shared support 1-Gbps, 2-Gbps, 4-Gbps in a maximum of 44 ports, or 8 Gbps in a maximum of 4 ports.

Dedicated Rate Mode Configurations for the 8-Gbps Modules

Table 3-7 shows the maximum possible dedicated rate mode configuration scenarios for the Generation 4 Fibre Channel modules.

Table 3-7 Dedicated Rate Mode Bandwidth Reservation for Generation 4 Fibre Channel Modules

Part Number	Product Name/ Description	Dedicated Bandwidth per Port	Maximum Allowed Ports That Can Come Up	Ports in Shared Mode
DS-X9248-256K9	48-port 8-Gbps Advanced Fibre Channel switching module	10 Gbps	24 Ports	All the remaining ports are 8 Gbps shared.
		8 Gbps	32 Ports	
		4 Gbps	48 Ports	
		2 Gbps	48 Ports	
		1 Gbps	48 Ports	
DS-X9232-256K9	32-port 8-Gbps Advanced Fibre Channel switching module	10 Gbps	24 Ports	All the remaining ports are 8 Gbps shared.
		8 Gbps	32 Ports	
		4 Gbps	32 Ports	
		2 Gbps	32 Ports	
		1 Gbps	32 Ports	

Table 3-8 shows the maximum possible dedicated rate mode configuration scenarios for the Generation 3 Fibre Channel modules.

Table 3-8 Dedicated Rate Mode Bandwidth Reservation for Generation 3 Fibre Channel Modules

Part Number	Product Name/ Description	Dedicated Bandwidth per Port	Maximum Allowed Ports That Can Come Up	Ports in Shared Mode
DS-X9224-96K9	24-port 8-Gbps Fibre Channel switching module	8 Gbps	8 Ports	All the remaining ports are 8 Gbps shared.
		4 Gbps	24 Ports	
DS-X9248-96K9	48-port 8-Gbps Fibre Channel switching module	8 Gbps	8 Ports	All the remaining ports are 8 Gbps shared.
		4 Gbps	24 Ports	
		2 Gbps	48 Ports	
DS-X9248-48K9	4/44-port 8-Gbps Host-Optimized Fibre Channel switching module	8 Gbps	4 Ports	All the remaining ports are 4 Gbps shared (8 Gbps of bandwidth can be provided only to one port per port group in dedicated or shared rate mode).
		4 Gbps	12 Ports	
		2 Gbps	24 Ports	
		1 Gbps	48 Ports	

Port Speed

The port speed on an interface, combined with the rate mode, determines the amount of shared resources available to the ports in the port group on a 48-port, 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module. Especially in the case of dedicated rate mode, the port group resources are reserved even though the bandwidth is not used. For example, on Generation 2 modules, if an interface is configured for autosensing (auto) and dedicated rate mode, then 4 Gbps of bandwidth is reserved even though the maximum operating speed is 2 Gbps. For the same interface, if autosensing with a maximum speed of 2 Gbps (auto max 2000) is configured, then only 2 Gbps of bandwidth is reserved and the unused 2 Gbps is shared with the other interface in the port group.

Note

- The Generation 2, 4-port 10-Gbps switching module supports 10-Gbps traffic only.
- On Generation 2, 4-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 4 Gbps.
- On Generation 3, 8-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 8 Gbps.
- On Generation 4, 8-Gbps modules, setting the port speed to auto enables autosensing, which negotiates to a maximum speed of 8 Gbps.

Dynamic Bandwidth Management

On port switching modules where bandwidth is shared, the bandwidth available to each port within a port group can be configured based on the port rate mode and speed configurations. Within a port group, some ports can be configured in dedicated rate mode while others operate in shared mode.

Ports configured in dedicated rate mode are allocated the required bandwidth to sustain a line rate of traffic at the maximum configured operating speed, and ports configured in shared mode share the available remaining bandwidth within the port group. Bandwidth allocation among the shared mode ports is based on the operational speed of the ports. For example, if four ports operating at speeds 1 Gbps, 1 Gbps, 2 Gbps, and 4 Gbps share bandwidth of 8 Gbps, the ratio of allocation would be 1:1:2:4.

Unutilized bandwidth from the dedicated ports is shared among only the shared ports in a port group as per the ratio of the configured operating speed. A port cannot be brought up unless the reserved bandwidth is guaranteed for the shared ports (see Table 3-9). For dedicated ports, configured bandwidth is taken into consideration while calculating available bandwidth for the port group. This behavior can be changed using bandwidth fairness by using the **rate-mode bandwidth-fairness module** *number* command.

For example, consider a 48-port 8-Gbps module. This module has 6 ports per port group with 12.8 Gbps bandwidth. Ports 3 to 6 are configured at 4 Gbps. If the first port is configured at 8 Gbps dedicated rate mode, and the second port is configured at 4-Gbps dedicated rate mode, then no other ports can be configured at 4 Gbps or 8 Gbps because the left over bandwidth of 0.8 Gbps (12.8-(8+4)) cannot meet the required 0.96 Gbps for the remaining four ports. A minimum of 0.24 Gbps reserved bandwidth is required for the for the rest of the four ports. However, if the two ports (for example, 5 and 6) are taken out of service (which is not same as shutdown), required reserved bandwidth for the two ports (3 and 4) is 0.48 and port 2 can be configured at 4 Gbps in dedicated rate mode. This behavior can be overridden by the bandwidth fairness command in which case reserved bandwidth is not enforced. Once the port is up, ports 3 and 4 can share the unutilized bandwidth from ports 1 and 2.

Out-of-Service Interfaces

On supported modules and fabric switches, you might need to allocate all the shared resources for one or more interfaces to another interface in the port group or module. You can take interfaces out of service to release shared resources that are needed for dedicated bandwidth. When an interface is taken out of service, all shared resources are released and made available to the other interface in the port group or module. These shared resources include bandwidth for the shared mode port, rate mode, BB_credits, and extended BB_credits. All shared resource configurations are returned to their default values when the interface is brought back into service. Corresponding resources must be made available in order for the port to be successfully returned to service.

Caution If you need to bring an interface back into service, you might disrupt traffic if you need to release shared resources from other interfaces in the same port group.

Oversubscription Ratio Restrictions

The 48-port and 24-port 4-Gbps, and all 8-Gbps Fibre Channel switching modules support oversubscription on switches with shared rate mode configurations. By default, all 48-port and 24-port 4-Gbps, and 8-Gbps Fibre Channel switching modules have restrictions on oversubscription ratios enabled. As of Cisco SAN-OS Release 3.1(1) and NX-OS Release 4.1(1), you can disable restrictions on oversubscription ratios.

Table 3-9 describes the bandwidth allocation for oversubscribed interfaces configured in shared mode on the 4-Gbps and 8-Gbps modules.

Table 3-9 Bandwidth Allocation for Oversubscribed Interfaces

Switching Module	Configured Speed	Reserved Bandwidth (Gbps)		Maximum Bandwidth (Gbps)
		Ratios enabled	Ratios disabled	
24-Port 8-Gbps Fibre Channel Module	Auto 8 Gbps	0.8	0.8	8
	Auto Max 4 Gbps	0.4	0.4	4
	Auto Max 2 Gbps	0.2	0.2	2
4/44-Port 8-Gbps Host-Optimized Fibre Channel Module	8 Gbps	0.87	0.16	8
	Auto Max 4 Gbps	0.436	0.08	4
	Auto Max 2 Gbps	0.218	0.04	2
	1 Gbps	0.109	0.02	1
48-port 4-Gbps Fibre Channel switching module	Auto 4 Gbps	0.8	0.09	4
	Auto Max 2 Gbps	0.4	0.045	2
	1 Gbps	0.2	0.0225	1
24-port 4-Gbps Fibre Channel switching module	Auto 4 Gbps	1	0.27	4
	Auto Max 2 Gbps	0.5	0.135	2
	1 Gbps	0.25	0.067	1

All ports in the 48-port and 24-port 4-Gbps modules can be configured to operate at 4 Gbps in shared mode even if other ports in the port group are configured in dedicated mode, regardless of available bandwidth. However, when oversubscription ratio restrictions are enabled, you may not have all shared 4-Gbps module ports operating at 4 Gbps.

All ports in the 48-port, 32-Port, and 24-port 8-Gbps modules can be configured to operate at 8 Gbps in shared mode even if other ports in the port group are configured in dedicated mode, regardless of available bandwidth. However, when oversubscription ratio restrictions are enabled you may not have all shared 8-Gbps module ports operating at 8 Gbps.

On the 48-port, 32-Port, and 24-port 8-Gbps modules, if you have configured one 8-Gbps dedicated port in one port group, no other ports in the same port group can be configured to operate at 8-Gbps dedicated mode. You can have any number of 8-Gbps shared and 4-Gbps dedicated or shared ports. On the 4/44-port 8-Gbps module, only one port per port group can be configured in 8-Gbps dedicated or shared mode.

In the following example, a 24-port 4-Gbps module has oversubscription ratios enabled and three dedicated ports in one port group operating at 4-Gbps. No other ports in the same port group can be configured to operate at 4 Gbps.

```
switch# show port-resources module 8
Module 8
  Available dedicated buffers are 5478

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 0.8 Gbps
  Allocated dedicated bandwidth is 12.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
```

```
  fc8/1                                          16        4.0  dedicated
  fc8/2                                          16        4.0  dedicated
  fc8/3                                          16        4.0  dedicated
  fc8/4 (out-of-service)
  fc8/5 (out-of-service)
  fc8/6 (out-of-service)
```

For dedicated ports, oversubscription ratio restrictions do not apply to the shared pool in port groups. So if oversubscription ratio restrictions are disabled, and you have configured three 4-Gbps dedicated ports in one port group, then you can configure all other ports in the same port group to operate at a shared rate of 4 Gbps.

In the following example, a 48-port module has a group of six ports, four dedicated ports are operating at 8 Gbps, and the two shared ports are also operating at 8 Gbps:

```
switch# show port-resources module 5
Module 5
Available dedicated buffers for global buffer #0 [port-group 1] are 3970
  Available dedicated buffers for global buffer #1 [port-group 2] are 3970
  Available dedicated buffers for global buffer #2 [port-group 3] are 3970
  Available dedicated buffers for global buffer #3 [port-group 4] are 3970
  Available dedicated buffers for global buffer #4 [port-group 5] are 3058
  Available dedicated buffers for global buffer #5 [port-group 6] are 3058
  Available dedicated buffers for global buffer #6 [port-group 7] are 3970
  Available dedicated buffers for global buffer #7 [port-group 8] are 3970

Port-Group 1
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/1                                      32        8.0  shared
  fc5/2                                      32        8.0  shared
  fc5/3                                      32        8.0  shared
  fc5/4                                      32        8.0  shared

Port-Group 2
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/5                                      32        8.0  shared
  fc5/6                                      32        8.0  shared
  fc5/7                                      32        8.0  shared
  fc5/8                                      32        8.0  shared

Port-Group 3
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/9                                      32        8.0  shared
  fc5/10                                     32        8.0  shared
  fc5/11                                     32        8.0  shared
```

```
  fc5/12                                    32        8.0  shared

Port-Group 4
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/13                                    32        8.0  shared
  fc5/14                                    32        8.0  shared
  fc5/15                                    32        8.0  shared
  fc5/16                                    32        8.0  shared

Port-Group 5
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 16.4 Gbps
  Allocated dedicated bandwidth is 16.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/17                                    32        8.0  shared
  fc5/18                                    32        8.0  shared
  fc5/19                                   500        8.0  dedicated
  fc5/20                                   500        8.0  dedicated

Port-Group 6
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 16.4 Gbps
  Allocated dedicated bandwidth is 16.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/21                                   500        8.0  dedicated
  fc5/22                                   500        8.0  dedicated
  fc5/23                                    32        8.0  shared
  fc5/24                                    32        8.0  shared

Port-Group 7
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/25                                    32        8.0  shared
  fc5/26                                    32        8.0  shared
  fc5/27                                    32        8.0  shared
  fc5/28                                    32        8.0  shared

Port-Group 8
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc5/29                                    32        8.0  shared
  fc5/30                                    32        8.0  shared
```

```
  fc5/31                                    32        8.0  shared
  fc5/32                                    32        8.0  shared


Isola-13# show port-resources module 13
Module 13
  Available dedicated buffers for global buffer #0 [port-group 1] are 3880
  Available dedicated buffers for global buffer #1 [port-group 2] are 3880
  Available dedicated buffers for global buffer #2 [port-group 3] are 3880
  Available dedicated buffers for global buffer #3 [port-group 4] are 3056
  Available dedicated buffers for global buffer #4 [port-group 5] are 3880
  Available dedicated buffers for global buffer #5 [port-group 6] are 3880
  Available dedicated buffers for global buffer #6 [port-group 7] are 3880
  Available dedicated buffers for global buffer #7 [port-group 8] are 3880

Port-Group 1
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/1                                    32        8.0  shared
  fc13/2                                    32        8.0  shared
  fc13/3                                    32        8.0  shared
  fc13/4                                    32        8.0  shared
  fc13/5                                    32        8.0  shared
  fc13/6                                    32        8.0  shared

Port-Group 2
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/7                                    32        8.0  shared
  fc13/8                                    32        8.0  shared
  fc13/9                                    32        8.0  shared
  fc13/10                                   32        8.0  shared
  fc13/11                                   32        8.0  shared
  fc13/12                                   32        8.0  shared

Port-Group 3
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/13                                   32        8.0  shared
  fc13/14                                   32        8.0  shared
  fc13/15                                   32        8.0  shared
  fc13/16                                   32        8.0  shared
  fc13/17                                   32        8.0  shared
  fc13/18                                   32        8.0  shared

Port-Group 4
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 0.4 Gbps
  Allocated dedicated bandwidth is 32.0 Gbps
```

```
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/19                                  250        8.0  dedicated
  fc13/20                                  250        8.0  dedicated
  fc13/21                                  250        8.0  dedicated
  fc13/22                                  250        8.0  dedicated
  fc13/23                                   32        8.0  shared
  fc13/24                                   32        8.0  shared

Port-Group 5
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/25                                   32        8.0  shared
  fc13/26                                   32        8.0  shared
  fc13/27                                   32        8.0  shared
  fc13/28                                   32        8.0  shared
  fc13/29                                   32        8.0  shared
  fc13/30                                   32        8.0  shared

Port-Group 6
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/31                                   32        8.0  shared
  fc13/32                                   32        8.0  shared
  fc13/33                                   32        8.0  shared
  fc13/34                                   32        8.0  shared
  fc13/35                                   32        8.0  shared
  fc13/36                                   32        8.0  shared

Port-Group 7
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc13/37                                   32        8.0  shared
  fc13/38                                   32        8.0  shared
  fc13/39                                   32        8.0  shared
  fc13/40                                   32        8.0  shared
  fc13/41                                   32        8.0  shared
  fc13/42                                   32        8.0  shared

Port-Group 8
  Total bandwidth is 32.4 Gbps
  Total shared bandwidth is 32.4 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
```

```
  fc13/43                                 32        8.0  shared
  fc13/44                                 32        8.0  shared
  fc13/45                                 32        8.0  shared
  fc13/46                                 32        8.0  shared
  fc13/47                                 32        8.0  shared
  fc13/48                                 32        8.0  shared

...
```

When disabling restrictions on oversubscription ratios, all ports in shared mode on 48-port and 24-port 4-Gbps or any 8-Gbps Fibre Channel switching modules must be shut down. When applying restrictions on oversubscription ratios, you must take shared ports out of service.

Note When restrictions on oversubscription ratios are disabled, the bandwidth allocation among the shared ports is proportionate to the configured speed. If the configured speed is auto on Generation 2 modules, then bandwidth is allocated assuming a speed of 4 Gbps. For example, if you have three shared ports configured at 1, 2, and 4 Gbps, then the allocated bandwidth ratio is 1:2:4.
As of Cisco SAN-OS Release 3.0 and NX-OS Release 4.1(1) or when restrictions on oversubscription ratios are enabled, the port bandwidths are allocated in equal proportions, regardless of port speed, so, the bandwidth allocation for the same three ports mentioned in the example would be 1:1:1.

Bandwidth Fairness

This feature improves fairness of bandwidth allocation among all ports and provides better throughput average to individual data streams. Bandwidth fairness can be configured per module.

As of Cisco SAN-OS Release 3.1(2), all 48-port and 24-port 4-Gbps Fibre Channel switching modules, as well as 18-port Fibre Channel/4-port Gigabit Ethernet Multiservice modules, have bandwidth fairness enabled by default. As of Cisco NX-OS Release 4.1(1), all the 8-Gbps Fibre Channel switching modules have bandwidth fairness enabled by default.

Caution When you disable or enable bandwidth fairness, the change does not take effect until you reload the module.

Use the **show module bandwidth-fairness** command to check whether ports in a module are operating with bandwidth fairness enabled or disabled.

```
switch# show module 2 bandwidth-fairness
Module 2 bandwidth-fairness is enabled
```

Note This feature is supported only on the 48-port and 24-port 4-Gbps modules, the 8-Gbps modules, and the 18/4-port Multiservice Module (MSM).

Upgrade or Downgrade Scenario

When you are upgrading from a release earlier than Cisco SAN-OS Release 3.1(2), all modules operate with bandwidth fairness disabled until the next module reload. After the upgrade, any new module that is inserted has bandwidth fairness enabled.

When you are downgrading to a release earlier than Cisco SAN-OS Release 3.1(2), all modules keep operating in the same bandwidth fairness configuration prior to the downgrade. After the downgrade, any new module that is inserted has bandwidth fairness disabled.

Note After the downgrade, any insertion of a module or module reload will have bandwidth fairness disabled.

Guidelines and Limitations

This section includes the following topics:

- Combining Generation 1, Generation 2, Generation 3, and Generation 4 Modules, page 3-18
- Local Switching Limitations, page 3-19
- Port Index Limitations, page 3-19
- PortChannel Limitations, page 3-21

Combining Generation 1, Generation 2, Generation 3, and Generation 4 Modules

Cisco MDS NX-OS Release 5.2(x) and later supports combining Generation 1, Generation 2, Generation 3, and Generation 4 modules and switches with the following considerations:

- MDS NX-OS Release 4.1(1) and later features are not supported on the following Generation 1 switches and modules:
  - Supervisor 1 module
  - 4-Port IP Storage Services module
  - 8-Port IP Storage Services module
  - MDS 9216 switch
  - MDS 9216A switch
  - MDS 9020 switch
  - MDS 9120 switch
  - MDS 9140 switch
- Supervisor-1 modules must be upgraded to Supervisor-2 modules on the MDS 9506 and MDS 9509 Directors.
- IPS-4 and IPS-8 modules must be upgraded to the MSM-18/4 Multiservice modules.
- Fabric 1 modules must be upgraded to Fabric 2 modules on the MDS 9513 Director to use the 48-port or the 24-port 8-Gbps module.
- Fabric 2 modules must be upgraded to Fabric 3 modules on the MDS 9513 Director to get the maximum backplane bandwidth of 256 Gbps.
- Cisco Fabric Manager Release 4.x supports MDS SAN-OS Release 3.x and NX-OS 4.x in mixed mode through Interswitch Link (ISL) connectivity.

Note When a Cisco or another vendor switch port is connected to a Generation 1 module port (ISL connection), the receive buffer-to-buffer credits of the port connected to the Generation 1 module port must not exceed 255.

Local Switching Limitations

- All ports in the module must be in shared mode. Use the **switchport ratemode shared** command to ensure that all the ports in the module are in shared mode.
- No E ports are allowed in the module because E ports must be in dedicated mode.

Port Index Limitations

Cisco MDS 9000 switches allocate index identifiers for the ports on the modules. These port indexes cannot be configured. You can combine Generation 1, Generation 2, Generation 3, and Generation 4 switching modules, with either Supervisor-1 modules or Supervisor-2 modules. However, combining switching modules and supervisor modules has the following port index limitations:

- Supervisor-1 modules only support a maximum of 252 port indexes, regardless of the type of switching modules.
- Supervisor-2 modules support a maximum of 1020 port indexes when all switching modules in the chassis are Generation 2 or Generation 3.
- Supervisor-2 modules only support a maximum of 252 port indexes when only Generation 1 switching modules, or a combination of Generation 1, Generation 2, Generation 3, or Generation 4 switching modules are installed in the chassis.

Note On a switch with the maximum limit of 252 as port index, any new module that exceeds the limit does not power up when installed.

You can use the **show port index-allocation** command to display the allocation of port indexes on the switch.

```
switch# show port index-allocation

Module index distribution:
-------------------------------------------------------+
Slot| Allowed |      Allotted indices info            |
    |  range  | Total |          Index values         |
----|---------|-------|-------------------------------|
1   |  -----  |   -   | (None)                        |
2   |  -----  |   -   | (None)                        |
3   |  -----  |   -   | (None)                        |
4   |  -----  |   -   | (None)                        |
5   |   0-1023|   32  | 0-31                          |
6   |  -----  |   -   | (None)                        |
9   |  -----  |   -   | (None)                        |
10  |  -----  |   -   | (None)                        |
11  |  -----  |   -   | (None)                        |
12  |  -----  |   -   | (None)                        |
13  |   0-1023|   48  | 32-79                         |
SUP | 253-255 |   3   | 253-255                       |
```

Generation 1 switching modules have specific numbering requirements. If these requirements are not met, the module does not power up. The port index numbering requirements include the following:

- If port indexes in the range of 256 to 1020 are assigned to operational ports, Generation 1 switching modules do not power up.
- A block of contiguous port indexes is available. If this block of port indexes is not available, Generation 1 modules do not power up. Table 3-10 shows the port index requirements for the Generation 1 modules.

Note If the switch has Supervisor-1 modules, the block of 32 contiguous port indexes must begin on the slot boundary. The slot boundary for slot 1 is 0, for slot 2 is 32, and so on. For Supervisor-2 modules, the contiguous block can start anywhere.

Table 3-10 Port Index Requirements for Generation 1 Modules

	Number of Port Indexes Required	
Generation 1 Module	**Supervisor-1 Module**	**Supervisor-2 Module**
16-port 2-Gbps Fibre Channel module	16	16
32-port 2-Gbps Fibre Channel module	32	32
8-port Gigabit Ethernet IP Storage Services module	32	32
4-port Gigabit Ethernet IP Storage Services module	32	16
32-port 2-Gbps Fibre Channel Storage Services Module (SSM).	32	32
14-port Fibre Channel/2-port Gigabit Ethernet Multiprotocol Services (MPS-14/2) module	32	22

The allowed mix of Generation 1 and Generation 2 switching modules in a chassis is determined at run-time, either when booting up the switch or when installing the modules. In some cases, the sequence in which switching modules are inserted into the chassis determines if one or more modules is powered up.

When a module does not power up because of a resource limitation, you can display the reason by using the **show module** command.

```
switch# show module
Mod  Ports  Module-Type                        Model              Status
---  -----  --------------------------------- ------------------ ----------
5    32     1/2/4/8/10 Gbps Advanced FC Module  DS-X9232-256K9     ok
7    0      Supervisor/Fabric-2                 DS-X9530-SF2-K9    active *
13   48     1/2/4/8/10 Gbps Advanced FC Module  DS-X9248-256K9     ok

Mod  Sw              Hw      World-Wide-Name(s) (WWN)
---  --------------  ------  --------------------------------------------------
5    5.2(2)          0.207   21:01:00:0d:ec:b7:28:c0 to 21:20:00:0d:ec:b7:28:c0
7    5.2(2)          1.9     --
13   5.2(2)          0.212   23:01:00:0d:ec:b7:28:c0 to 23:30:00:0d:ec:b7:28:c0


Mod  MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
5    68-ef-bd-a8-45-cc to 68-ef-bd-a8-45-d0  JAF1450CHQT
7    00-24-c4-60-00-f8 to 00-24-c4-60-00-fc  JAE141502L2
```

```
13   68-ef-bd-a8-40-00 to 68-ef-bd-a8-40-04  JAF1450BMBP

Xbar Ports  Module-Type                          Model              Status
---  -----  ----------------------------------- ------------------ ----------
1    0      Fabric Module 3                     DS-13SLT-FAB3      ok
2    0      Fabric Module 3                     DS-13SLT-FAB3      ok

Xbar Sw              Hw      World-Wide-Name(s) (WWN)
---  --------------  ------  --------------------------------------------------
1    NA              0.4     --
2    NA              0.4     --


Xbar MAC-Address(es)                         Serial-Num
---  --------------------------------------  ----------
1    NA                                      JAF1451AMHG
2    NA                                      JAF1451AMHN


* this terminal session
```

The running configuration is updated when modules are installed. If you save the running configuration to the startup configuration (using the **copy running-config startup-config** command), during reboot the switch powers up the same set of modules as before the reboot regardless of the sequence in which the modules initialize. You can use the **show port index-allocation startup** command to display the index allocation the switch uses at startup.

```
switch# show port index-allocation startup

Startup module index distribution:
-------------------------------------------------------+
Slot | Allowed |      Alloted indices info             |
     |  range  | Total |          Index values         |
-----|---------|-------|-------------------------------|
 1   |  -----  |   34  | 0-31,80-81                    |
 2   |  -----  |   32  | 32-63                         |
 3   |  -----  |   16  | 64-79                         |(Slot 1 shares 80-81)
 4   |  -----  |   48  | 96-127,224-239                |
SUP  | 253-255 |   3   | 253-255                       |
```

Note The output of the **show port index-allocation startup** command does not display anything in the Allowed range column because the command extracts the indices from the persistent storage service (PSS) and displaying an allowed range for startup indices is meaningless.

If a module fails to power up, you can use the **show module** *slot* **recovery-steps** command to display the reason.

For information on recovering a module powered-down because port indexes are not available, refer to the *Cisco MDS 9000 Family Troubleshooting Guide*.

Tip Whenever using mixed Generation 1 and Generation 2 modules, power up the Generation 1 modules first. During a reboot of the entire switch, the Generation 1 modules power up first (default behavior).

PortChannel Limitations

PortChannels have the following restrictions:

- The maximum number of PortChannels allowed is 256 if all switching modules are Generation 2 or Generation 3, or both.
- The maximum number of PortChannels allowed is 128 whenever there is a Generation 1 switching module in use with a Generation 2 or Generation 3 switching module.
- Ports need to be configured in dedicated rate mode on the Generation 2 and Generation 3 switching module interfaces to be used in the PortChannel.

Note The number of PortChannels allowed does not depend on the type of supervisor module. However, Generation 3 modules require the Supervisor 2 module on the MDS 9506 and 9509 switches.

The Generation1, Generation 2, and Generation 3 modules have the following restrictions for PortChannel configuration:

- Generation 1 switching module interfaces do not support auto speed with a maximum of 2 Gbps.
- Generation 1 and Generation 2 module interfaces do not support auto speed with maximum of 4 Gbps.
- Generation 2 and Generation 3 switching module interfaces cannot be forcefully added to a PortChannel if sufficient resources are not available.

Note Before adding a Generation 2 or Generation 3 interface to a PortChannel, use the **show port-resources module** command to check for resource availability.

When configuring PortChannels on switches with Generation 1, Generation 2, and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 2 and Generation 3 interfaces to auto with a maximum of 2 Gbps.
- Configure the Generation 1 switching modules followed by the Generation 2 switching modules, and then the Generation 3 switching modules, and then configure the PortChannel.

When configuring PortChannels on switches with only Generation 2 and Generation 3 switching modules, follow one of these procedures:

- Configure the PortChannel, and then configure the Generation 3 interfaces to auto with a maximum of 4 Gbps.
- Configure the Generation 2 switching modules, followed by the Generation 3 switching modules, and then configure the PortChannel.

Table 3-11 describes the results of adding a member to a PortChannel for various configurations.

Table 3-11 PortChannel Configuration and Addition Results

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
No members	Any	Any	Generation 1 or Generation 2 or Generation 3 or Generation 4	Force	Pass
	Auto	Auto	Generation 1 or Generation 2 or Generation 3 or Generation 4	Normal or force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass or fail[1]
	Auto	Auto max 4000	Generation 3 or Generation 4		
	Auto max 2000	Auto	Generation 2 or Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto max 2000	Auto max 4000	Generation 3 or or Generation 4		
	Auto max 4000	Auto	Generation 2 or Generation 3 or or Generation 4		
	Auto max 4000	Auto max 2000	Generation 2 or Generation 3 or or Generation 4		
Generation 1 interfaces	Auto	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass or fail[1]
	Auto max 4000	Auto	Generation 1 or Generation 2		
	Auto max 4000	Auto	Generation 3		

Table 3-11 PortChannel Configuration and Addition Results (continued)

PortChannel Members	Configured Speed		New Member Type	Addition Type	Result
	PortChannel	New Member			
Generation 2 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2 or Generation 3	Normal	Fail
				Force	Pass
Generation 3 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3	Normal	Fail
				Force	Pass
Generation 4 interfaces	Auto	Auto	Generation 1	Normal or force	Fail
	Auto max 2000	Auto	Generation 1	Normal or force	Pass
	Auto max 2000	Auto	Generation 2	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 2	Normal	Fail
				Force	Pass
	Auto max 2000	Auto	Generation 3 or Generation 4	Normal	Fail
				Force	Pass
	Auto	Auto max 2000	Generation 3 or Generation 4	Normal	Fail
				Force	Pass

1. If resources are not available.

Use the **show port-channel compatibility parameters** command to obtain information about PortChannel addition errors.

Default Settings

Table 3-12 lists the default settings for Generation 2 interface parameters.

Table 3-12 Default Generation 2 Interface Parameters

Parameter	Default			
	48-Port 4-Gbps Switching Module	24-Port 4-Gbps Switching Module	12-Port 4-Gbps Switching Module	4-Port 10-Gbps Switching Module
Speed mode	**auto**	**auto**[1]	**auto**[1]	**auto**[1]
Rate mode	**shared**	**shared**	**dedicated**	**dedicated**
Port mode	Fx	Fx	auto[2]	auto[3]
BB_credit buffers	16	16	250	250
Performance buffers	–	–	145[4]	145[5]

1. The 4-port 10-Gbps switching module only supports 10-Gbps traffic.
2. Auto port mode on the 12-port 4-Gbps switching module interfaces can operate in E port mode, TE port mode, and Fx port mode.
3. Auto port mode on the 4-port 10-Gbps switching module interfaces can operate in E port mode, TE port mode, and F port mode.
4. Performance buffers are shared among all ports on the module.

Table 3-13 lists the default settings for Generation 3 interface parameters.

Table 3-13 Default Generation 3 Interface Parameters

Parameter	Default		
	48-Port 8-Gbps Switching Module	24-Port 8-Gbps Switching Module	4/44-Port 8-Gbps Host-Optimized Switching Module
Speed mode	**auto**	**auto**[1]	**auto_max_4G**[1]
Rate mode	**shared**	**shared**	**shared**
Port mode	Fx	Fx	Fx
BB_credit buffers	32	32	32

1. Auto_max_4G speed mode on the 4/44-port 8-Gbps switching module negotiates to a maximum speed of 4 Gbps.

Table 3-14 lists the default settings for Generation 4 interface parameters.

Table 3-14 Default Generation 4 Interface Parameters

Parameter	Default	
	48-Port 8-Gbps Advanced Fibre Channel Switching Module	32-Port 8-Gbps Advanced Fibre Channel Switching Module
Speed mode	**auto**	**auto**[1]
Rate mode	**shared**	**shared**

Table 3-14 Default Generation 4 Interface Parameters (continued)

Parameter	Default	
	48-Port 8-Gbps Advanced Fibre Channel Switching Module	32-Port 8-Gbps Advanced Fibre Channel Switching Module
Port mode	Fx	Fx
BB_credit buffers	32	32

Configuring Fibre Channel Interfaces

This section includes the following topics:

- Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode, page 3-26
- Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode, page 3-27
- Task Flow for Configuring 12-Port 4-Gbps Module Interfaces, page 3-28
- Task Flow for Configuring 4-Port 10-Gbps Module Interfaces, page 3-28
- Configuring Port Speed, page 3-29
- Configuring Rate Mode, page 3-30
- Configuring Local Switching, page 3-36
- Disabling Restrictions on Oversubscription Ratios, page 3-37
- Enabling Restrictions on Oversubscription Ratios, page 3-39
- Enabling Bandwidth Fairness, page 3-40
- Disabling Bandwidth Fairness, page 3-41
- Taking Interfaces Out of Service, page 3-41
- Releasing Shared Resources in a Port Group, page 3-42
- Disabling ACL Adjacency Sharing for System Image Downgrade, page 3-43

Task Flow for Migrating Interfaces from Shared Mode to Dedicated Mode

The 48-Port, 24-Port, and 4/44-Port 8-Gbps Fibre Channel switching modules support the following features:

- 1-Gbps, 2-Gbps, 4-Gbps, and 8-Gbps speed traffic
- Shared and dedicated rate mode
- ISL and Fx port modes
- Extended BB_credits

The 48-port and 24-port 4-Gbps Fibre Channel switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Shared and dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits

Note If you change the port bandwidth reservation parameters on a 48-port or 24-port 4-Gbps module, the change affects only the changed port. No other ports in the port group are affected.

To configure the 4-Gbps and 8-Gbps Fibre Channel switching modules when starting with the default configuration or when migrating from shared rate mode to dedicated rate mode, follow these steps:

Step 1 Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the "Taking Interfaces Out of Service" section on page 3-41.

Step 2 Configure the traffic speed to use (1 Gbps, 2 Gbps, 4 Gbps, 8 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps).

See the "Configuring Port Speed" section on page 3-29.

Step 3 Configure the rate mode (dedicated or shared).

See the "Configuring Rate Mode" section on page 3-30.

Step 4 Configure the port mode.

See the "Configuring Interface Modes" section on page 2-22.

Note ISL ports cannot operate in shared rate mode.

Step 5 Configure the BB_credits and extended BB_credits, as necessary.

See the "Configuring Buffer-to-Buffer Credits" section on page 4-21 and the "Configuring Extended BB_credits" section on page 4-23.

Task Flow for Migrating Interfaces from Dedicated Mode to Shared Mode

To configure the 4-Gbps and 8-Gbps Fibre Channel switching modules migrating from dedicated rate mode to shared rate mode, follow these steps:

Step 1 Take unused interfaces out of service to release resources for other interfaces, if necessary.

See the "Taking Interfaces Out of Service" section on page 3-41.

Step 2 Configure the BB_credits and extended BB_credits, as necessary.

See the "Configuring Buffer-to-Buffer Credits" section on page 4-21, and the "Extended BB_Credits" section on page 4-17.

Step 3 Configure the port mode.

See the "Configuring Interface Modes" section on page 2-22.

Note ISL ports cannot operate in shared rate mode.

Step 4 Configure the rate mode (dedicated or shared) to use.

See the "Configuring Rate Mode" section on page 3-30.

Step 5 Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

See the "Configuring Port Speed" section on page 3-29.

Task Flow for Configuring 12-Port 4-Gbps Module Interfaces

The 12-port 4-Gbps switching modules support the following features:

- 1-Gbps, 2-Gbps, and 4-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and Fx (F or FL) port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these steps:

Step 1 Configure the traffic speed (1 Gbps, 2 Gbps, 4 Gbps, or autosensing with a maximum of 2 Gbps or 4 Gbps) to use.

See the "Configuring Port Speed" section on page 3-29.

Step 2 Configure the port mode.

See the "Configuring Interface Modes" section on page 2-22.

Step 3 Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

See the "Configuring Buffer-to-Buffer Credits" section on page 4-21, and the "Configuring Extended BB_credits" section on page 4-23.

Task Flow for Configuring 4-Port 10-Gbps Module Interfaces

The 4-port 10-Gbps switching modules support the following features:

- Only 10-Gbps speed traffic
- Only dedicated rate mode
- ISL (E or TE) and F port modes
- Extended BB_credits
- Performance buffers

To configure 4-port 10-Gbps switching modules when starting with the default configuration, follow these steps:

Step 1 Configure the port mode.

See the "Configuring Interface Modes" section on page 2-22.

Step 2 Configure the BB_credits, performance buffers, and extended BB_credits, as necessary.

See the "Configuring Buffer-to-Buffer Credits" section on page 4-21, and the "Configuring Extended BB_credits" section on page 4-23.

Configuring Port Speed

Restrictions

- Changing port speed and rate mode disrupts traffic on the port. Traffic on other ports in the port group is not affected.

Detailed Steps

To configure the port speed on an interface on a 4-Gbps or 8-Gbps switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport speed 4000`	Configures the port speed in megabits per second. Valid values are 1000, 2000, 4000 and auto. The auto parameter enables autosensing on the interface.
	`switch(config-if)# switchport speed 8000`	(8-Gbps modules only[1]) Configures the port speed in megabits per second to 8-Gbps. Valid values are 1000, 2000, 4000, 8000 and auto. The auto parameter enables autosensing on the interface.
	`switch(config-if)# switchport speed auto`	On 4-Gbps modules, configures autosensing for the interface with 4 Gbps of bandwidth reserved. On 8-Gbps modules, configures autosensing for the interface with 8 Gbps of bandwidth reserved. [2]
	`switch(config-if)# switchport speed auto max 2000`	Configures autosensing with a maximum of 2 Gbps of bandwidth reserved.
	`switch(config-if)# switchport speed auto max 4000`	(8-Gbps modules only[1]) Configures autosensing with a maximum of 4Gbps of bandwidth reserved.
	`switch(config-if)# no switchport speed`	Reverts to the default[2] speed for the interface (auto).

1. The 8000 and auto max 4000 speed configurations are available only for the 8-Gbps modules.
2. The default speed on 48-port and 24-port 4-Gbps modules is 4000. The default speed on 48-port and 24-port 8-Gbps modules is 8000. The default speed on the 4/44-port 8-Gbps module is auto max 4000.

Use the **show interface** command to verify the port speed configuration for an interface on a 4-Gbps or 8-Gbps Fibre Channel switching module.

```
switch# show interface fc 9/1
fc9/1 is up
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 22:01:00:05:30:01:9f:02
    Admin port mode is F
    snmp traps are enabled
    Port mode is F, FCID is 0xeb0002
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is shared
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      226 frames input, 18276 bytes
        0 discards, 0 errors
        0 CRC, 0 unknown class
        0 too long, 0 too short
      326 frames output, 21364 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 1 NOS, 0 loop inits
      3 output OLS, 2 LRR, 0 NOS, 0 loop inits
      16 receive B2B credit remaining
      64 transmit B2B credit remaining
```

Configuring Rate Mode

Restrictions

- Changing port speed and rate mode disrupts traffic on the port.

Detailed Steps

To configure the rate mode (dedicated or shared) on an interface on a 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport rate-mode dedicated`	Reserves dedicated bandwidth for the interface. **Note** If you cannot reserve dedicated bandwidth on an interface, you might have exceeded the port group maximum bandwidth. Use the **show port-resources** command to determine what resources are already allocated.
	`switch(config-if)# switchport rate-mode shared`	Reserves shared (default) bandwidth for the interface.
	`switch(config-if)# no switchport rate-mode`	Reverts to the default state (**shared**).

Displaying the Rate Mode Configuration for Interfaces

Use **show port-resources module** command to verify the rate mode configuration for interfaces on a 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module.

This example shows the port rate mode configuration for interfaces on a 4-Gbps module:

```
switch# show port-resources module 9
Module 9
Available dedicated buffers are 5400

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/1                                     16        4.0  shared
  fc9/2                                     16        4.0  shared
  fc9/3                                     16        4.0  shared
  fc9/4                                     16        4.0  shared
  fc9/5                                     16        4.0  shared
  fc9/6                                     16        4.0  shared

 Port-Group 2
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------
  fc9/7                                     16        4.0  shared
  fc9/8                                     16        4.0  shared
  fc9/9                                     16        4.0  shared
  fc9/10                                    16        4.0  shared
  fc9/11                                    16        4.0  shared
  fc9/12                                    16        4.0  shared

 Port-Group 3
  Total bandwidth is 12.8 Gbps
```

```
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------
  fc9/13                                    16        4.0  shared
  fc9/14                                    16        4.0  shared
  fc9/15                                    16        4.0  shared
  fc9/16                                    16        4.0  shared
  fc9/17                                    16        4.0  shared
  fc9/18                                    16        4.0  shared

 Port-Group 4
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------
  fc9/19                                    16        4.0  shared
  fc9/20                                    16        4.0  shared
  fc9/21                                    16        4.0  shared
  fc9/22                                    16        4.0  shared
  fc9/23                                    16        4.0  shared
  fc9/24                                    16        4.0  shared
```

This example shows the port rate mode configuration for interfaces on a 48-port 8-Gbps module:

```
switch# show port-resource module 4
Module 4
  Available dedicated buffers for global buffer #0 [port-groups 1-4] are 5016
  Available dedicated buffers for global buffer #1 [port-groups 5-8] are 5016

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------
  fc4/1                                     32        8.0  shared
  fc4/2                                     32        8.0  shared
  fc4/3                                     32        8.0  shared
  fc4/4                                     32        8.0  shared
  fc4/5                                     32        8.0  shared
  fc4/6                                     32        8.0  shared

 Port-Group 2
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------
  fc4/7                                     32        8.0  shared
  fc4/8                                     32        8.0  shared
  fc4/9                                     32        8.0  shared
  fc4/10                                    32        8.0  shared
  fc4/11                                    32        8.0  shared
```

```
 fc4/12                                    32        8.0  shared

Port-Group 3
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc4/13                                    32        8.0  shared
 fc4/14                                    32        8.0  shared
 fc4/15                                    32        8.0  shared
 fc4/16                                    32        8.0  shared
 fc4/17                                    32        8.0  shared
 fc4/18                                    32        8.0  shared

Port-Group 4
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc4/19                                    32        8.0  shared
 fc4/20                                    32        8.0  shared
 fc4/21                                    32        8.0  shared
 fc4/22                                    32        8.0  shared
 fc4/23                                    32        8.0  shared
 fc4/24                                    32        8.0  shared

Port-Group 5
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc4/25                                    32        8.0  shared
 fc4/26                                    32        8.0  shared
 fc4/27                                    32        8.0  shared
 fc4/28                                    32        8.0  shared
 fc4/29                                    32        8.0  shared
 fc4/30                                    32        8.0  shared

Port-Group 6
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc4/31                                    32        8.0  shared
 fc4/32                                    32        8.0  shared
 fc4/33                                    32        8.0  shared
 fc4/34                                    32        8.0  shared
 fc4/35                                    32        8.0  shared
 fc4/36                                    32        8.0  shared

Port-Group 7
 Total bandwidth is 12.8 Gbps
```

```
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------------
  fc4/37                                    32        8.0  shared
  fc4/38                                    32        8.0  shared
  fc4/39                                    32        8.0  shared
  fc4/40                                    32        8.0  shared
  fc4/41                                    32        8.0  shared
  fc4/42                                    32        8.0  shared

 Port-Group 8
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------------
  fc4/43                                    32        8.0  shared
  fc4/44                                    32        8.0  shared
  fc4/45                                    32        8.0  shared
  fc4/46                                    32        8.0  shared
  fc4/47                                    32        8.0  shared
  fc4/48                                    32        8.0  shared
```

This example shows the port rate mode configuration for interfaces on a 4/44-port 8-Gbps module:

```
switch# show port-resources module 7
Module 7
  Available dedicated buffers are 3888

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------------
  fc7/1                                     32        4.0  shared
  fc7/2                                     32        4.0  shared
  fc7/3                                     32        4.0  shared
  fc7/4                                     32        4.0  shared
  fc7/5                                     32        4.0  shared
  fc7/6                                     32        4.0  shared
  fc7/7                                     32        4.0  shared
  fc7/8                                     32        4.0  shared
  fc7/9                                     32        4.0  shared
  fc7/10                                    32        4.0  shared
  fc7/11                                    32        4.0  shared
  fc7/12                                    32        4.0  shared

 Port-Group 2
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  ------------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  ------------------------------------------------------------------------
```

```
 fc7/13                                        32        4.0  shared
 fc7/14                                        32        4.0  shared
 fc7/15                                        32        4.0  shared
 fc7/16                                        32        4.0  shared
 fc7/17                                        32        4.0  shared
 fc7/18                                        32        4.0  shared
 fc7/19                                        32        4.0  shared
 fc7/20                                        32        4.0  shared
 fc7/21                                        32        4.0  shared
 fc7/22                                        32        4.0  shared
 fc7/23                                        32        4.0  shared
 fc7/24                                        32        4.0  shared

Port-Group 3
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc7/25                                        32        4.0  shared
 fc7/26                                        32        4.0  shared
 fc7/27                                        32        4.0  shared
 fc7/28                                        32        4.0  shared
 fc7/29                                        32        4.0  shared
 fc7/30                                        32        4.0  shared
 fc7/31                                        32        4.0  shared
 fc7/32                                        32        4.0  shared
 fc7/33                                        32        4.0  shared
 fc7/34                                        32        4.0  shared
 fc7/35                                        32        4.0  shared
 fc7/36                                        32        4.0  shared

Port-Group 4
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc7/37                                        32        4.0  shared
 fc7/38                                        32        4.0  shared
 fc7/39                                        32        4.0  shared
 fc7/40                                        32        4.0  shared
 fc7/41                                        32        4.0  shared
 fc7/42                                        32        4.0  shared
 fc7/43                                        32        4.0  shared
 fc7/44                                        32        4.0  shared
 fc7/45                                        32        4.0  shared
 fc7/46                                        32        4.0  shared
 fc7/47                                        32        4.0  shared
 fc7/48                                        32        4.0  shared
```

Configuring Local Switching

Note We recommend that you shut down all of the ports on the module before you execute the local switching command.

If local switching is enabled, then ports cannot be configured in dedicated mode.

If there are dedicated ports and you enter the local switching command, a warning is displayed and the operation is prevented.

Detailed Steps

To enable or disable local switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# local-switching module 5`	Enables local switching for a module. If you want to proceed with enabling local switching, when prompted press **Y** to continue or **N** to exit.
Step 3	`switch(config)# no local-switching module 5`	Disables local switching for a module. If you want to proceed with disabling local switching, when prompted press **Y** to continue or **N** to exit.
Step 4	`switch(config)# exit`	Exits configuration mode.

Use the **show system internal xbar local-switching** command to verify the local switching configuration status on a module.

```
switch# show system internal xbar local-switching
-------------------------------------
| Slot | Local-switching enabled  |
-------------------------------------
|   1  |           no             |
|   2  |           no             |
|   3  |           no             |
|   4  |           no             |
|   5  |           yes            |
|   6  |           no             |
|   9  |           no             |
|  10  |           no             |
|  11  |           no             |
|  12  |           no             |
|  13  |           no             |
-------------------------------------
```

Disabling Restrictions on Oversubscription Ratios

Prerequisites

- Before disabling restrictions on oversubscription ratios, ensure that you have explicitly shut down shared ports.

Detailed Steps

To disable restrictions on oversubscription ratios on a 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no rate-mode oversubscription-limit module 1`	Disables restrictions on oversubscription ratios for a module. **Note** You must enter this command separately for each module for which you want to remove the restrictions.
Step 3	`switch(config)# exit`	Exits configuration mode.
Step 4	`switch# copy running-config startup-config`	Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Use the **show running-config** command to view oversubscription ratios for a module. If oversubscription ratios are enabled, then no restriction appears in the output.

Example 3-1 Module with Restrictions on Oversubscription Ratios Disabled

```
switch# show running-config
version 3.1(1)
...
no rate-mode oversubscription-limit module 2
interface fc2/1
  switchport speed 2000
interface fc2/1
...
```

Examples

To disable restrictions on oversubscription ratios for ports on a 48-port Generation 2 switch that is configured with both shared and dedicated ports, follow these steps:

Step 1 To disable restrictions on oversubscription ratios, you must shut down any shared ports. Use the **show port-resources** command to view the configuration on a module and to identify shared ports.

```
switch# show port-resources module 2
Module 2
Available dedicated buffers are 4656
```

```
Port-Group 1
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc2/1                                      16        4.0  shared
 fc2/2                                      16        4.0  shared
 fc2/3                                      16        4.0  dedicated
 fc2/4                                      16        4.0  shared
 fc2/5                                      16        4.0  shared
 fc2/6                                      16        4.0  dedicated
 fc2/7                                      16        4.0  dedicated
 fc2/8                                      16        4.0  shared
 fc2/9                                      16        4.0  shared
 fc2/10                                     16        4.0  shared
 fc2/11                                     16        4.0  shared
 fc2/12                                     16        4.0  shared
 ...

Port-Group 4
 Total bandwidth is 12.8 Gbps
 Total shared bandwidth is 12.8 Gbps
 Allocated dedicated bandwidth is 0.0 Gbps
 --------------------------------------------------------------------
 Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                       Buffers     (Gbps)
 --------------------------------------------------------------------
 fc2/37                                     16        4.0  shared
 fc2/38                                     16        4.0  shared
 fc2/39                                     16        4.0  dedicated
 fc2/40                                     16        4.0  dedicated
 fc2/41                                     16        4.0  dedicated
 fc2/42                                     16        4.0  shared
 fc2/43                                     16        4.0  shared
 fc2/44                                     16        4.0  shared
 fc2/45                                     16        4.0  shared
 fc2/46                                     16        4.0  shared
 fc2/47                                     16        4.0  shared
 fc2/48                                     16        4.0  shared
```

Step 2 Shut down all shared ports for which you want to remove restrictions on oversubscription ratios.

```
switch (config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch (config-if)# shutdown
```

Step 3 Display the interface status to confirm the shutdown of all shared ports.

```
switch(config-if)# end
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                 Mode   Trunk                         Mode  Speed  Channel
                        Mode                                (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      down             swl    --           --
fc2/2      1      FX     --      down             swl    --           --
fc2/3      1      T      --      up               swl    --           --
fc2/4      1      FX     --      down             swl    --           --
fc2/5      1      FX     --      down             swl    --           --
fc2/6      1      TE     --      up               swl    --           --
```

```
fc2/7      1      TE     --      up              swl    --          --
fc2/8      1      FX     --      down            swl    --          --
...
fc2/48     1      FX     --      down            sw1    --          --
```

Step 4 Disable restrictions on oversubscription ratios for the ports.

```
switch# config t
Enter configuration commands, one per line.  End with CNTL/Z.
switch(config)# no rate-mode oversubscription-limit module 2
```

Step 5 Bring up the ports that you shut down in step 2, and display their status to confirm that they are no longer shut down.

```
switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48
switch(config-if)# no shutdown
switch(config-if)# end
switch# show interface brief

--------------------------------------------------------------------------------
Interface  Vsan   Admin  Admin   Status          SFP    Oper  Oper   Port
                  Mode   Trunk                         Mode  Speed  Channel
                         Mode                                (Gbps)
--------------------------------------------------------------------------------
fc2/1      1      FX     --      up              swl    --          --
fc2/2      1      FX     --      up              swl    --          --
fc2/3      1      T      --      up              swl    --          --
fc2/4      1      FX     --      up              swl    --          --
fc2/5      1      FX     --      up              swl    --          --
fc2/6      1      TE     --      up              swl    --          --
fc2/7      1      TE     --      up              swl    --          --
fc2/8      1      FX     --      up              swl    --          --
...
fc2/48     1      FX     --      up              sw1    --          --
```

Step 6 Confirm that the ports are now operating with no restrictions on oversubscription ratios.

```
switch# show running-config | include oversubscription-limit
no rate-mode oversubscription-limit module 2 <---indicates no restrictions on
                                              oversubscription ratios
```

Step 7 Save the new oversubscription ratio configuration to the startup configuration.

```
switch# copy running-config startup-config
```

Enabling Restrictions on Oversubscription Ratios

Prerequisites

- You must enable restrictions on oversubscription ratios before you can downgrade modules to a previous release.
- Before enabling restrictions on oversubscription ratios, ensure that you have explicitly configured shared ports to out-of-service mode.

Detailed Steps

To enable restrictions on oversubscription ratios on a 48-port or 24-port 4-Gbps, or any 8-Gbps Fibre Channel switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48`	Specifies the port interfaces for which you want to enable restrictions on oversubscription ratios.
Step 3	`switch(config-if)# shutdown`	Shuts down shared ports.
Step 4	`switch(config-if)# out-of-service`	Takes shared ports out of service.
Step 5	`switch# rate-mode oversubscription-limit module 1`	Enables restrictions on oversubscription ratios for the module. **Note** You must enter this command separately for each module for which you want to add the restriction.

	`switch# config t`
	`switch(config)# interface fc2/1-2, fc2/4-5, fc2/8-38, fc2/43-48`
	`switch(config-if)# no out-of-service`
	`switch(config-if)# no shutdown`
		Returns all shared ports to service.

Step 6	`switch(config)# exit`	Exits configuration mode.
Step 7	`switch# copy running-config startup-config`	Saves the new oversubscription ratio configuration to the startup configuration, and then the new configuration is enforced upon subsequent reboots of the module.

Enabling Bandwidth Fairness

Detailed Steps

To enable bandwidth fairness on a switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# rate-mode bandwidth-fairness module 1`	Enables bandwidth fairness for a module. **Note** You must enter this command separately for each module for which you want to enable bandwidth fairness. You must reload the module for the command to take effect.
Step 3	`switch(config)# exit`	Exits configuration mode.

Disabling Bandwidth Fairness

Restrictions

- If you disable bandwidth fairness, up to a 20 percent increase in internal bandwidth allocation is possible for each port group; however, bandwidth fairness is not guaranteed when there is a mix of shared and full-rate ports in the same port group.

Detailed Steps

To disable bandwidth fairness on a switching module, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# no rate-mode bandwidth-fairness module 1`	Disables bandwidth fairness for a module. **Note** You must enter this command separately for each module for which you want to disable bandwidth fairness. You must reload the module for the command to take effect.
Step 3	`switch(config)# exit`	Exits configuration mode.

Taking Interfaces Out of Service

You can take interfaces out of service on Generation 2 and Generation 3 switching modules. When an interface is out of service, all the shared resources for the interface are released as well as the configuration associated with those resources.

Prerequisites

- The interface must be disabled using a **shutdown** command before it can be taken out of service.

Restrictions

- The interface cannot be a member of a PortChannel.
- Taking interfaces out of service releases all the shared resources to ensure that they are available to other interfaces. This causes the configuration in the shared resources to revert to default when the interface is brought back into service. Also, an interface cannot come back into service unless the default shared resources for the port are available. The operation to free up shared resources from another port is disruptive.

Detailed Steps

To take an interface out of service, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# no channel-group`	Removes the interface from a PortChannel.
Step 4	`switch(config-if)# shutdown`	Disables the interface.

Step 5	`switch(config-if)# out-of-service`
	`Putting an interface into out-of-service will cause its shared resource configuration to revert to default`
	`Do you wish to continue (y/n)? [n] y`
		Takes the interface out of service.

Use the **show port-resources module** command to verify the out-of-service configuration for interfaces on a Generation 2 and Generation 3 switching module.

This example shows a 24-port 4-Gbps module:

```
switch# show port-resources module 9
Module 9
Available dedicated buffers are 5429

 Port-Group 1
  Total bandwidth is 12.8 Gbps
  Total shared bandwidth is 12.8 Gbps
  Allocated dedicated bandwidth is 0.0 Gbps
  --------------------------------------------------------------------------
  Interfaces in the Port-Group       B2B Credit  Bandwidth  Rate Mode
                                        Buffers     (Gbps)
  --------------------------------------------------------------------------
  fc9/1                                     16        4.0  shared
  fc9/2 (out-of-service)
  fc9/3                                     16        4.0  shared
  fc9/4                                     16        4.0  shared
  fc9/5                                     16        4.0  shared
  fc9/6                                     16        4.0  shared
...
```

Releasing Shared Resources in a Port Group

When you want to reconfigure the interfaces in a port group on a Generation 2 or Generation 3 module, you can return the port group to the default configuration to avoid problems with allocating shared resources.

Restrictions

- The interface cannot be a member of a PortChannel.
- Releasing shared resources disrupts traffic on the port. Traffic on other ports in the port group is not affected.

Detailed Steps

To release the shared resources for a port group, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode. **Tip** You can use an interface range to release the resources for all interfaces in a port group.
Step 3	`switch(config-if)# no channel-group`	Removes the interface from a PortChannel.
Step 4	`switch(config-if)# shutdown`	Disables the interface.

Step 5	`switch(config-if)# out-of-service`
	`Putting an interface into out-of-service will cause its shared resource configuration to revert to default`
	`Do you wish to continue(y/n)? [n] y`
		Takes the interface out of service.

Step 6	`switch(config-if)# no out-of-service`	Makes the interface available for service. Repeat Step 2 through Step 6 for all the interfaces in the port group.

Disabling ACL Adjacency Sharing for System Image Downgrade

Fibre Channel ACL adjacency sharing is enabled by default on the switches with an active Generation 2 switching module as of Cisco MDS SAN-OS Release 3.0(3), and with an active Generation 3 module as of MDS NX-OS Release 4.1(1). Fibre Channel ACL adjacency sharing improves the performance for zoning and inter-VSAN routing (IVR) network address translation (NAT). To prevent disruptions when downgrading the system image on your switch to a release prior to Cisco SAN-OS Release 3.0(3), enter the following command in EXEC mode:

```
switch# system no acl-adjacency-sharing
```

To reenable Fibre Channel ACL adjacency sharing on your switch, enter the following command in EXEC mode:

```
switch# system acl-adjacency-sharing
```

Verifying Fibre Channel Interfaces Configuration

To display Fibre Channel interface configuration information, perform one of the following tasks:

Command	Purpose
show module	Displays the module.
show module *slot* **recovery-steps**	Displays the slot for the module.
show port-resources module *slot*	Displays the port resources for the slot.

Command	Purpose
show interface fc *slot/port*	Displays the slot or port information.
show interface brief	Displays the interface.
show port index-allocation	Displays the port in the index allocation.
show port index-allocation startup	Displays the startup port in the index allocation.
show port-channel compatibility parameters	Displays the PortChannel compatibility parameters.
show module *slot* **bandwidth-fairness**	Displays the module slot bandwidth fairness information.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

Displaying Interface Capabilities

Before configuring a Generation 2 or Generation 3 interface, you can use the **show interface capabilities** command to display detailed information about the capabilities of the interface.

This example shows the capabilities of a Generation 2 Fibre Channel interface:

```
switch# show interface fc 9/1 capabilities
 Min Speed is 1 Gbps
 Max Speed is 4 Gbps
 FC-PH Version (high, low)                        (0,6)
 Receive data field size (max/min)                (2112/256) bytes
 Transmit data field size (max/min)               (2112/128)  bytes
 Classes of Service supported are                 Class 2, Class 3, Class F
 Class 2 sequential delivery                      supported
 Class 3 sequential delivery                      supported
 Hold time (max/min)                              (100/1) micro sec
 BB state change notification                     supported
 Maximum BB state change notifications            14
 Rate Mode change                                 supported

 Rate Mode Capabilities                           Shared        Dedicated
  Receive BB Credit modification supported           yes             yes
  FX mode Receive BB Credit (min/max/default)     (1/16/16)      (1/250/16)
  ISL mode Receive BB Credit (min/max/default)        --         (2/250/250)
  Performance buffer modification supported           no              no

 Out of Service capable                            yes
 Beacon mode configurable                          yes
```

This example shows the capabilities of an interface on the 48-port 8-Gbps Fibre Channel switching module:

```
switch# show interface fc 4/1 capabilities
 Min Speed is 1 Gbps
 Max Speed is 8 Gbps
 FC-PH Version (high, low)                        (0,6)
 Receive data field size (max/min)                (2112/256) bytes
 Transmit data field size (max/min)               (2112/128) bytes
 Classes of Service supported are                 Class 2, Class 3, Class F
 Class 2 sequential delivery                      supported
 Class 3 sequential delivery                      supported
 Hold time (max/min)                              (100000/1) micro sec
```

```
 BB state change notification                      supported
 Maximum BB state change notifications             14
 Rate Mode change                                  supported

 Rate Mode Capabilities                            Shared        Dedicated
  Receive BB Credit modification supported           yes             yes
  FX mode Receive BB Credit (min/max/default)     (1/32/32)      (1/500/32)
  ISL mode Receive BB Credit (min/max/default)        --         (2/500/250)
  Performance buffer modification supported           no              no

 Out of Service capable                            yes
 Beacon mode configurable                          yes
 Extended B2B credit capable                       yes
```

Displaying SFP Diagnostic Information

You can use the **show interface transceiver** command to display small form-factor pluggable (SFP) diagnostic information for Generation 2 switching modules.

```
switch# show interface transceiver
...
fc12/12 sfp is present
    name is CISCO-FINISAR
    part number is FTRJ-8519-7D2CS1
    revision is A
    serial number is H11TVQB
    fc-transmitter type is short wave laser w/o OFC (SN)
    fc-transmitter supports intermediate distance link length
    media type is multi-mode, 62.5m (M6)
    Supported speed is 200 MBytes/sec
    Nominal bit rate is 2100 MBits/sec
    Link length supported for 50/125mm fiber is 300 m(s)
    Link length supported for 62.5/125mm fiber is 150 m(s)
    cisco extended id is unknown (0x0)

    no tx fault, rx loss, no sync exists, Diag mon type 104
    SFP Diagnostics Information
        Temperature      :  24.33 Celsius
        Voltage          :   3.33 Volt
        Current          :   0.04 mA       --
        Optical Tx Power :  N/A   dBm      --
        Optical Rx Power :  N/A   dBm      -
    Note: ++  high-alarm; +  high-warning; --  low-alarm; -  low-warning
...
```

Configuration Examples for Fibre Channel Interfaces

This section describes example configurations and includes the following sections:

- Configuration Example for 48-Port 8-Gbps Module Interfaces, page 3-46
- Configuration Example for 24-Port 8-Gbps Module Interfaces, page 3-46
- Configuration Example for 4/44-Port 8-Gbps Module Interfaces, page 3-47
- Configuration Example for 48-Port 4-Gbps Module Interfaces, page 3-48
- Configuration Example for 24-Port 4-Gbps Module Interfaces, page 3-49

Configuration Example for 48-Port 8-Gbps Module Interfaces

These steps describe how to configure the 48-port 8-Gbps module interfaces:

Step 1 Select the interfaces fc 4/1 through fc 4/2.

```
switch# config t
switch(config)# interface fc 4/1 - 2
```

Step 2 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 8000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 3 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 4/3 through fc 4/4.

```
switch# config t
switch(config)# interface fc 4/3 - 4
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed auto max 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode f
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 7 Select the interfaces fc 4/5 through fc 4/6.

```
switch# config t
switch(config)# interface fc 4/5 - 6
```

Step 8 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed auto max 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 9 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuration Example for 24-Port 8-Gbps Module Interfaces

These steps describe how to configure the 24-port 8-Gbps module interfaces:

Step 1 Select interfaces fc 3/1.

```
switch# config t
switch(config)# interface fc 3/1
```

Step 2 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 8000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode f
```

Step 3 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 3/2 through fc 3/3.

```
switch# config t
switch(config)# interface fc 3/2 - 3
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 8000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuration Example for 4/44-Port 8-Gbps Module Interfaces

These steps describe how to configure the 4/44-port 8-Gbps module interfaces:

Step 1 Select interfaces fc 4/11 through fc 4/12.

```
switch# config t
switch(config)# interface fc 4/11 - 12
```

Step 2 Disable the interfaces and take them out of service.

```
switch(config-if)# shutdown
switch(config-if)# out-of-service
```

Step 3 Return to configuration mode.

```
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 4/1.

```
switch# config t
switch(config)# interface fc 4/1
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 8000
```

```
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 7 Select the interfaces fc 4/2 through fc 4/10.

```
switch# config t
switch(config)# interface fc 4/2 - 10
```

Step 8 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed auto max 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 9 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuration Example for 48-Port 4-Gbps Module Interfaces

These steps describe how to configure the example shown in Figure 4-6 on page 4-11:

Step 1 Select interfaces fc 4/11 through fc 4/12.

```
switch# config t
switch(config)# interface fc 4/11 - 12
```

Step 2 Disable the interfaces and take them out of service.

```
switch(config-if)# shutdown
switch(config-if)# out-of-service
```

Step 3 Return to configuration mode.

```
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 4/1 through fc 4/6.

```
switch# config t
switch(config)# interface fc 4/1 - 6
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed auto max 2000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode e
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 7 Select the interfaces fc 4/7 through fc 4/10.

```
switch# config t
switch(config)# interface fc 4/7 - 10
```

Step 8 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 1000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 9 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuration Example for 24-Port 4-Gbps Module Interfaces

These steps describe how to configure the example shown in Figure 4-8 on page 4-13:

Step 1 Select interfaces fc 3/1 through fc 3/3.

```
switch# config t
switch(config)# interface fc 3/1 - 3
```

Step 2 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport mode e
```

Step 3 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Step 4 Select the interfaces fc 3/4 through fc 3/6.

```
switch# config t
switch(config)# interface fc 3/4 - 6
```

Step 5 Configure the port speed, rate mode, and port mode on the interfaces.

```
switch(config-if)# switchport speed 1000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport mode f
```

Step 6 Enable the interfaces and return to configuration mode.

```
switch(config-if)# no shutdown
switch(config-if)# exit
switch#
```

Configuring Interface Buffers

- Information About Interface Buffers, page 4-1
- Configuring Interface Buffers, page 4-21
- Verifying BB_Credit Configuration, page 4-25

Information About Interface Buffers

Fibre Channel interfaces use buffer credits to ensure all packets are delivered to their destination.

This section includes the following topics:

- Buffer-to-Buffer Credits, page 4-1
- Performance Buffers, page 4-2
- Buffer Pools, page 4-2
- BB_Credit Buffers for Switching Modules, page 4-5
- BB_Credit Buffers for Fabric Switches, page 4-16
- Extended BB_Credits, page 4-17
- Buffer-to-Buffer Credit Recovery, page 4-20
- Buffer-to-Buffer State Change Number, page 4-20
- Receive Data Field Size, page 4-21

Buffer-to-Buffer Credits

Buffer-to-buffer credits (BB_credits) are a flow-control mechanism to ensure that Fibre Channel switches do not run out of buffers, so that switches do not drop frames. BB_credits are negotiated on a per-hop basis.

The receive BB_credit (fcrxbbcredit) value may be configured for each Fibre Channel interface. In most cases, you do not need to modify the default configuration.

The receive BB_credit values depend on the module type and the port mode, as follows:

- For 16-port switching modules and full rate ports, the default value is 16 for Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required.
- For 32-port switching modules and host-optimized ports, the default value is 12 for Fx, E, and TE modes. These values cannot be changed.

- For Generation 2, Generation 3, and Generation 4 switching modules, see the "Buffer Pools" section on page 4-2.

Note In the Cisco MDS 9100 Series switches, the groups of ports on the left outlined in white are in dedicated rate mode. The other ports are host-optimized. Each group of 4 host-optimized ports have the same features as for the 32-port switching module.

Note Because Generation 1 modules do not support as many buffer-to-buffer credits as Generation 4 modules supports, you cannot configure an ISL on E or TE ports between a Generation 1 module such as the 16-port 1-, 2-Gbps Fibre Channel Switching Module (DS-X9016) and a Generation 4 module such as the 48 port 8-Gbps Advanced Fibre Channel module (DS-X9248-256K9) or the 32-port 8-Gbps Advanced Fibre Channel module (DS-X9232-256K9).

Performance Buffers

Regardless of the configured receive BB_credit value, additional buffers, called performance buffers, improve switch port performance. Instead of relying on the built-in switch algorithm, you can manually configure the performance buffer value for specific applications (for example, forwarding frames over FCIP interfaces).

Note Performance buffers are not supported on the Cisco MDS 9148 Fabric Switch, Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

For each physical Fibre Channel interface in any switch in the Cisco MDS 9000 Family, you can specify the amount of performance buffers allocated in addition to the configured receive BB_credit value.

The default performance buffer value is 0. If you set the performance buffer value to 0, the built-in algorithm is used. If you do not specify the performance buffer value, 0 is automatically used.

The default performance buffer value is 0. If you use the **default** option, the built-in algorithm is used. If you do not specify this command, the **default** option is automatically used.

Buffer Pools

In the architecture of Generation 2, Generation 3, and Generation 4 modules, receive buffers shared by a set of ports are called *buffer groups*. The receive buffer groups are organized into *global* and *local* buffer pools.

The receive buffers allocated from the global buffer pool to be shared by a port group are called a *global receive buffer pool*. Global receive buffer pools include the following buffer groups:

- Reserved internal buffers
- Allocated BB_credit buffers for each Fibre Channel interface (user configured or assigned by default)
- Common unallocated buffer pool for BB_credits, if any, to be used for additional BB_credits as needed
- Performance buffers (only used on 12-port 4-Gbps and 4-port 10-Gbps switching modules)

Note The 48-port and 24-port 8-Gbps modules have *dual global buffer pools*. Each buffer pool in the 48-port modules support 24 ports and in the 24-port modules each buffer pool supports 12 ports.

Figure 4-1 shows the allocation of BB_credit buffers on line cards (24-port and 48-port 4-Gbps line cards).

Figure 4-1 Receive Buffers for Fibre Channel Ports in a Global Buffer Pool

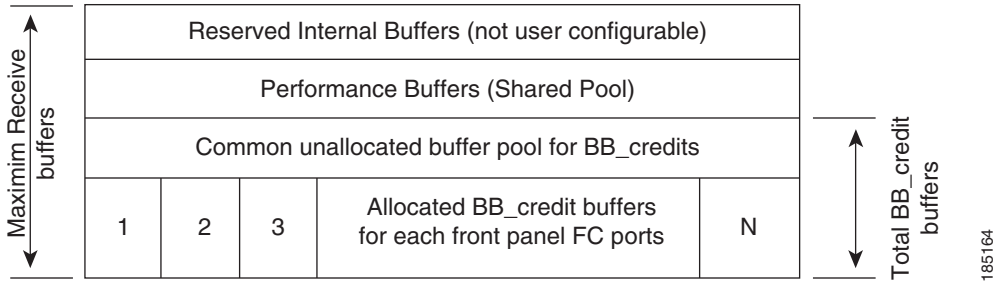

Figure 4-2 shows the default BB_credit buffer allocation model for 48-port 8-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 4-2 BB_Credit Buffer Allocation in 48-Port 8-Gbps Switching Modules

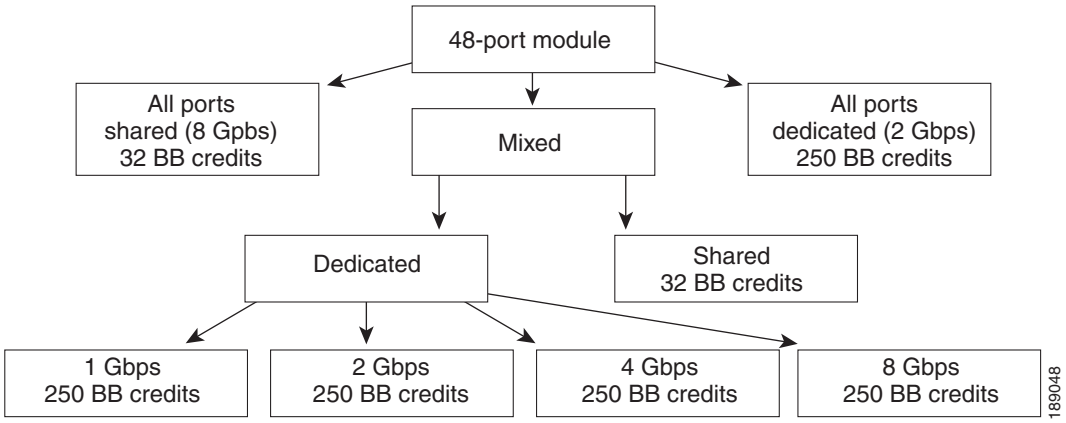
48-port module
All ports shared (8 Gpbs) 32 BB credits
Mixed
All ports dedicated (2 Gbps) 250 BB credits
Dedicated
Shared 32 BB credits
1 Gbps 250 BB credits
2 Gbps 250 BB credits
4 Gbps 250 BB credits
8 Gbps 250 BB credits
189048

Figure 4-3 shows the default BB_credit buffer allocation model for 24-port 8-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 4-3 BB_Credit Buffer Allocation in 24-Port 8-Gbps Switching Modules

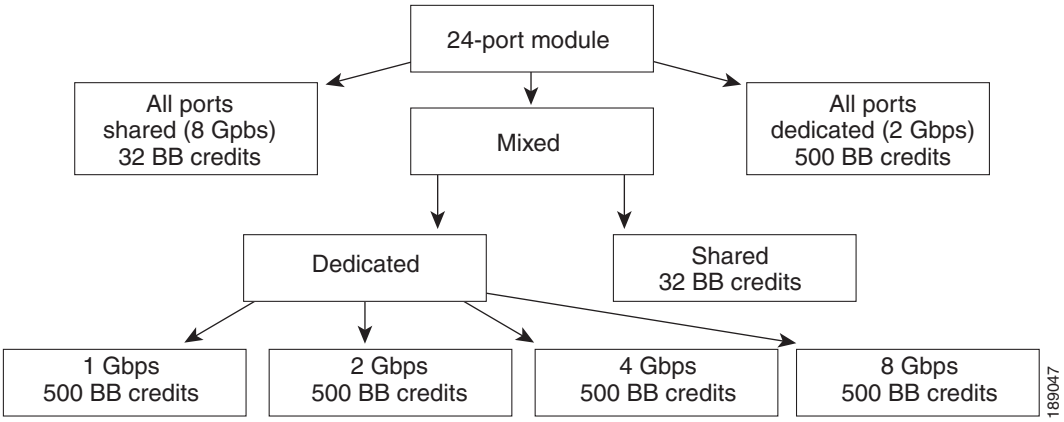

Figure 4-4 shows the default BB_credit buffer allocation model for 4/44-port 8-Gbps host-optimized switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 4-4 BB_Credit Buffer Allocation in 4/44-Port 8-Gbps Switching Modules

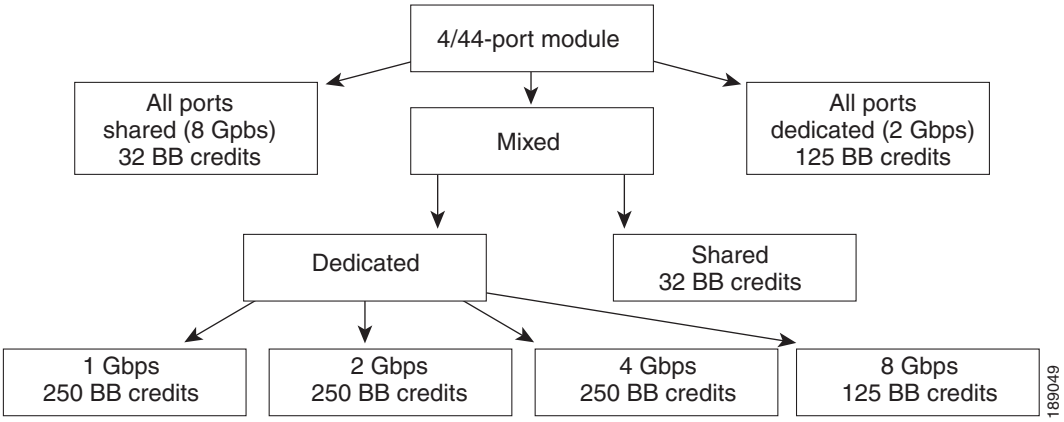

Figure 4-5 shows the default BB_credit buffer allocation model for 24-port 4-Gbps switching modules. The minimum BB_credits required to bring up a port is two buffers.

Figure 4-5 BB_Credit Buffer Allocation in 24-Port 4-Gbps Switching Modules

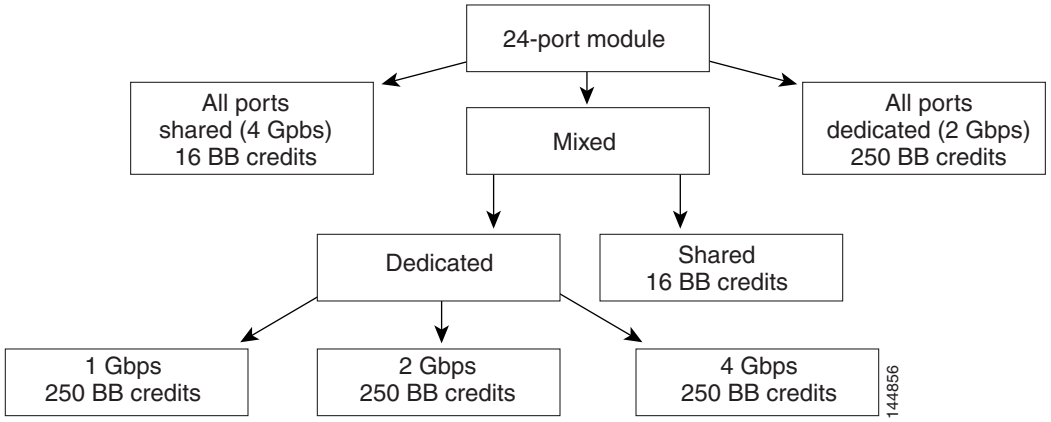

Note The default BB_credit buffer allocation is the same for all port speeds.

BB_Credit Buffers for Switching Modules

This section describes how buffer credits are allocated to Cisco MDS 9000 switching modules, and includes the following topics:

- Configuring Buffer Credits on a Generation 2, Generation 3 or Generation 4 Module, page 4-5
- 48-Port 8-Gbps Advanced Fibre Channel Module BB_Credit Buffers, page 4-6
- 48-Port 8-Gbps Fibre Channel Module BB_Credit Buffers, page 4-7
- 24-Port 8-Gbps Fibre Channel Module BB_Credit Buffers, page 4-8
- 4/44-Port 8-Gbps Host-Optimized Fibre Channel Module BB_Credit Buffers, page 4-9
- 48-Port 4-Gbps Fibre Channel Module BB_Credit Buffers, page 4-10
- 24-Port 4-Gbps Fibre Channel Module BB_Credit Buffers, page 4-12
- 18-Port Fibre Channel/4-Port Gigabit Ethernet Multiservice Module BB_Credit Buffers, page 4-13
- 4-Port 10-Gbps Switching Module BB_Credit Buffers, page 4-15

Configuring Buffer Credits on a Generation 2, Generation 3 or Generation 4 Module

When you configure port mode to auto or E on a Generation 2 module, one of the ports will not come up for the following configuration:

- Port Mode: auto or E for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E on a Generation 3 module, one or two of the ports will not come up for the following configuration:

- Port Mode: auto or E for the first half of the ports, the second half of the ports or for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E for all ports in the global buffer pool, you need to reconfigure buffer credits on one or more of the ports. The total number of buffer credits configured for all the ports in the global buffer pool should be reduced by 64.

48-Port 8-Gbps Advanced Fibre Channel Module BB_Credit Buffers

Table 4-1 lists the BB_credit buffer allocation for the 48-port 8-Gbps Advanced Fibre Channel switching module..

Table 4-1 48-Port 8-Gbps Advanced Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 8-Gbps Speed		Shared Rate Mode 8-Gbps Speed
BB_Credit Buffer Allocation	ISL	Fx Port	Fx Port
Default BB_credit buffers	250 for 48 port 500 for 32 port	32	32
Maximum BB_credit buffers	500	500	32

The following guidelines apply to BB_credit buffers on 32/48-port Advanced 8-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.
- The buffers should not be allocated automatically.

Each port group on the 32/48-port Advanced 8-Gbps Fibre Channel switching module consists of four/six ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 1.5:1 considering that each port group has 32-Gbps bandwidth. In case of 32 Port version, each port group of 4 ports has sufficient bandwidth (32 Gbps) to handle the line rate traffic without any oversubscription.

The following example configurations are supported by the 48-port Advanced 8-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 8-Gbps speed (1.5:1 oversubscription) (default).
- Two port with dedicated rate mode and 8-Gbps speed plus four ports with shared rate mode and 8-Gbps speed (2:1 oversubscription).

- Two ports with dedicated rate mode and 8-Gbps speed plus four ports with shared rate mode and 8-Gbps speed (2:1 oversubscription) .
- One port with dedicated rate mode and 8-Gbps speed plus three ports with dedicated rate mode and 4-Gbps speed plus two ports with shared rate mode and 8-Gbps speed (1.33:1 oversubscription).
- Six ports with dedicated rate mode and 8-Gbps speed.

48-Port 8-Gbps Fibre Channel Module BB_Credit Buffers

Table 4-2 lists the BB_credit buffer allocation for the 48-port 8-Gbps Fibre Channel switching module.

Table 4-2 48-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 8-Gbps Speed		Shared Rate Mode 8-Gbps Speed
BB_Credit Buffer Allocation	**ISL**	**Fx Port**	**Fx Port**
Default BB_credit buffers	250	32	32
Maximum BB_credit buffers	500	500	32
Total Number of BB_Credit Buffers per Module			
Ports 1 through 24	6000		
Ports 25 through 48	6000		

The following guidelines apply to BB_credit buffers on 48-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers allocated for ports 1 through 24 and 25 through 48 can be a maximum of 6000 each so that the load is distributed.
- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.
- The buffers should not be allocated automatically.

Each port group on the 48-port 8-Gbps Fibre Channel switching module consists of six ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 48-port 8-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 8-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
  five ports with shared rate mode and 8-Gbps speed (10:1 oversubscription)
- Two ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)

- One port with dedicated rate mode and 4-Gbps speed plus
  three ports with dedicated rate mode and 2-Gbps speed plus
  two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
- Six ports with dedicated rate mode and 2-Gbps speed

24-Port 8-Gbps Fibre Channel Module BB_Credit Buffers

Table 4-3 lists the BB_credit buffer allocation for the 24-port 8-Gbps Fibre Channel switching module.

Table 4-3 24-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 8-Gbps Speed		Shared Rate Mode 8-Gbps Speed
BB_Credit Buffer Allocation	**ISL**	**Fx Port**	**Fx Port**
Default BB_credit buffers	500	32	32
Maximum BB_credit buffers	500[1]	500[1]	32
Total Number of BB_Credit Buffers per Module			
Ports 1 through 12	6000		
Ports 13 through 24	6000		

1. When connected to Generation 1 modules, reduce the maximum BB_credit allocation to 250.

The following guidelines apply to BB_credit buffers on 24-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers allocated for ports 1 through 12 and 13 through 24 can be a maximum of 6000 each so that the load is distributed.
- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 500 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 8-Gbps Fibre Channel switching module consists of three ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 24-port 8-Gbps Fibre Channel switching modules:

- Three ports with shared rate mode and 8-Gbps speed (2:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
  two ports with shared rate mode and 8-Gbps speed (4:1 oversubscription)
- One port with dedicated rate mode and 8-Gbps speed plus
  one port with dedicated rate mode and 4-Gbps speed plus
  one port with shared rate mode and 8-Gbps speed (10:1 oversubscription)
- Two ports with dedicated rate mode and 4-Gbps speed plus
  one port with shared rate mode and 8-Gbps speed (2:1 oversubscription)

- Three ports with dedicated rate mode and 4-Gbps speed

4/44-Port 8-Gbps Host-Optimized Fibre Channel Module BB_Credit Buffers

Table 4-4 lists the BB_credit buffer allocation for the 4/44-port 8-Gbps Fibre Channel switching module.

Table 4-4 4/44-Port 8-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 8-Gbps Speed		Shared Rate Mode 8-Gbps Speed
BB_Credit Buffer Allocation	ISL	Fx Port	Fx Port
Default BB_credit buffers	125	32	32
Maximum BB_credit buffers	250	250	32
Total number of BB_credit buffers per module	6000		

The following guidelines apply to BB_credit buffers on 4/44-port 8-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 500 buffers for dedicated rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 32 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 8-Gbps Fibre Channel switching module consists of 12 ports. The ports in shared rate mode in a port group can have a maximum bandwidth oversubscription of 10:1 considering that each port group has 12.8-Gbps bandwidth.

The following example configurations are supported by the 4/44-port 8-Gbps Fibre Channel switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (5:1 oversubscription) (default)
- One port with dedicated rate mode and 8-Gbps speed plus
  eleven ports with shared rate mode and 4-Gbps speed (10:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  three ports with dedicated rate mode and 3-Gbps speed plus
  eight ports with shared rate mode and 4-Gbps speed (2:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed

48-Port 4-Gbps Fibre Channel Module BB_Credit Buffers

Table 4-5 lists the BB_credit buffer allocation for 48-port 4-Gbps Fibre Channel switching modules.

Table 4-5 48-Port 4-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed
BB_Credit Buffer Allocation	ISL[1]	Fx Port	Fx Port
Default BB_credit buffers	125	16	16
Maximum BB_credit buffers	250	250	16
Total number of BB_credit buffers per module	6000		

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 48-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 48-port 4-Gbps Fibre Channel switching module consists of 12 ports. The ports in shared rate mode have bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 48-port 4-Gbps Fibre Channel switching modules:

- Twelve ports with shared rate mode and 4-Gbps speed (4:1 oversubscription) (default)
- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  11 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 4-Gbps speed (5:1 oversubscription)
- Two ports with dedicated rate mode and 2-Gbps speed plus
  10 ports with shared rate mode and 2-Gbps speed (2.5:1 oversubscription)
- Twelve ports with dedicated rate mode and 1-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 1-Gbps speed plus
  five ports put out-of-service (see Figure 4-6)

Figure 4-6 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

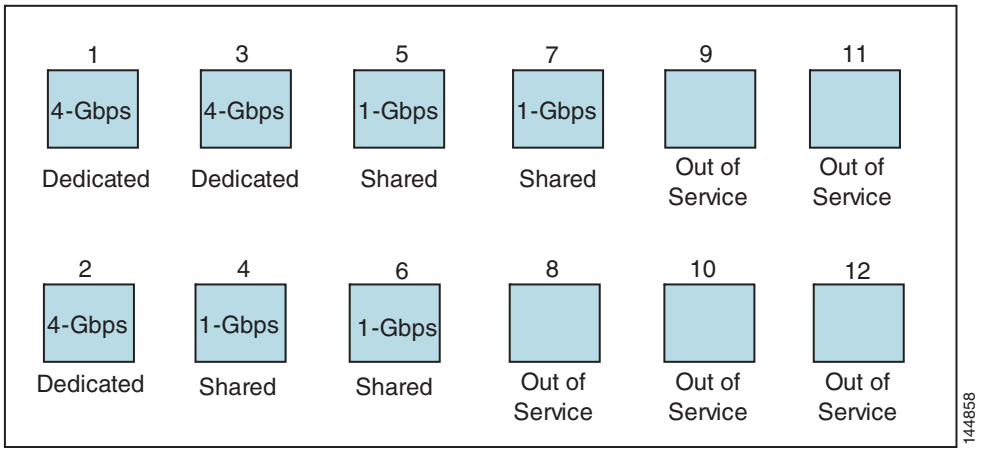

- Six ports with dedicated rate mode and 2-Gbps speed plus four ports with shared rate mode and 1-Gbps speed plus two ports put out-of-service (see Figure 4-7)

Note For detailed configuration steps of this example, see "Configuration Example for 48-Port 4-Gbps Module Interfaces" section on page 3-48.

Figure 4-7 Example Speed and Rate Configuration on a 48-Port 4-Gbps Switching Module

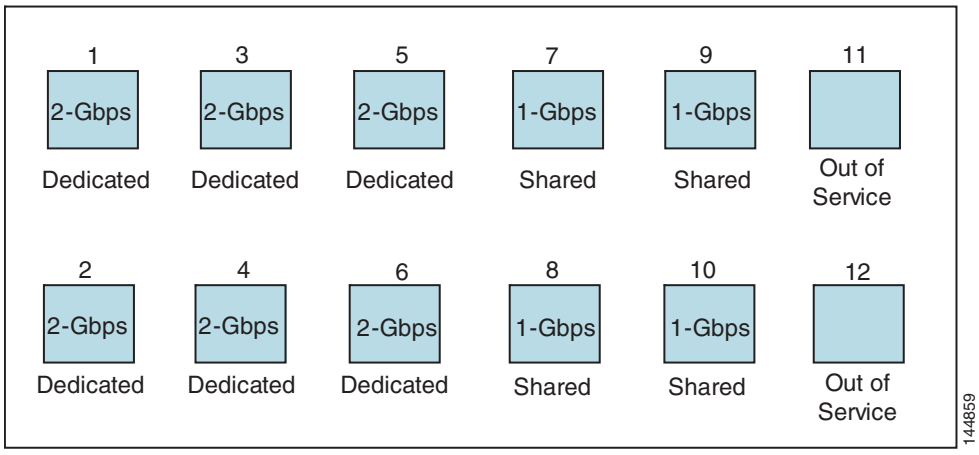

24-Port 4-Gbps Fibre Channel Module BB_Credit Buffers

Table 4-6 lists the BB_credit buffer allocation for 24-port 4-Gbps Fibre Channel switching modules.

Table 4-6 24-Port 4-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port		
	Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed
BB_Credit Buffer Allocation	ISL[1]	Fx Port	Fx Port
Default BB_credit buffers	250	16	16
Maximum BB_credit buffers	250	250	16
Total number of BB_credits buffers per module	6000		

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 24-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

Each port group on the 24-port 4-Gbps Fibre Channel switching module consists of six ports. The ports in shared rate mode have a bandwidth oversubscription of 2:1 by default. However, some configurations of the shared ports in a port group can have a maximum bandwidth oversubscription of 4:1 (considering that each port group has 12.8-Gbps bandwidth).

The following example configurations are supported by the 24-port 4-Gbps Fibre Channel switching modules:

- Six ports with shared rate mode and 4-Gbps speed (2:1 oversubscription) (default)
- Two ports with dedicated rate mode and 4-Gbps speed plus
  four ports with shared rate mode and 4-Gbps speed (with 4:1 oversubscription)
- One port with dedicated rate mode and 4-Gbps speed plus
  three ports with dedicated rate mode and 2-Gbps speed plus
  two ports with shared rate mode and 4-Gbps speed (4:1 oversubscription)
- Six ports with dedicated rate mode and 2-Gbps speed
- Three ports with dedicated rate mode and 4-Gbps speed plus
  three ports with shared rate mode and 1-Gbps speed (see Figure 4-8)

Note For detailed configuration steps of this example, see the "Configuration Example for 24-Port 4-Gbps Module Interfaces" section on page 3-49.

Figure 4-8 Example Speed and Rate Configuration on a 24-Port 4-Gbps Switching Module

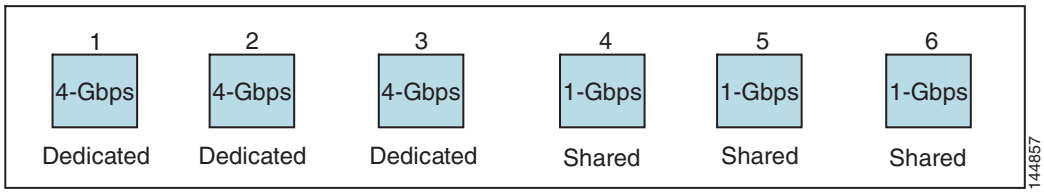

18-Port Fibre Channel/4-Port Gigabit Ethernet Multiservice Module BB_Credit Buffers

Table 4-7 lists the BB_credit buffer allocation for 18-port 4-Gbps multiservice modules.

Table 4-7 18-Port 4-Gbps Multiservice Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port			
	Dedicated Rate Mode 4-Gbps Speed		Shared Rate Mode 4-Gbps Speed	
BB_Credit Buffer Allocation	ISL[1]	Fx Port	ISL[1]	Fx Port
Default BB_credit buffers	250	16	16	16
Maximum BB_credit buffers	250	250	16	16
Total number of BB_credit buffers per module	4509			

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on18-port 4-Gbps Fibre Channel switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- BB_credit buffers for Fx port mode connections can be configured. The minimum is 2 buffers and the maximum of 250 buffers for dedicated rate mode or 16 buffers for shared rate mode.
- Performance buffers are not supported on this module.

12-Port 4-Gbps Switching Module BB_Credit Buffers

Table 4-8 lists the BB_credit buffer allocation for 12-port 4-Gbps switching modules.

Table 4-8 12-Port 4-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port	
	Dedicated Rate Mode 4-Gbps Speed	
BB_Credit Buffer Allocation Type	ISL [1]	Fx Port
Default BB_credit buffers	250	16
Maximum BB_credit buffers	250	16
Default Performance buffers	145	12

Table 4-8 12-Port 4-Gbps Switching Module BB_Credit Buffer Allocation (continued)

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port	
	Dedicated Rate Mode 4-Gbps Speed	
	ISL [1]	Fx Port
Total number of BB_credit buffers per module	5488	
Total number of performance buffers per module	512 (shared)	

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 12-port 4-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 250 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credit buffers after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (250 * 12)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

Note By default, the ports in the 12-port 4-Gbps switching modules come up in 4-Gbps dedicated rate mode but can be configured as 1-Gbps and 2-Gbps dedicated rate mode. Shared mode is not supported.

4-Port 10-Gbps Switching Module BB_Credit Buffers

Table 4-9 lists the BB_credit buffer allocation for 4-port 10-Gbps switching modules.

Table 4-9 4-Port 10-Gbps Switching Module BB_Credit Buffer Allocation

	BB_Credit Buffers Per Port	
	Dedicated Rate Mode 10-Gbps Speed	
BB_Credit Buffer Allocation Type	ISL[1]	F port[2]
Default BB_credit buffers	250	16
Maximum BB_credit buffers	750	16
Maximum BB_credit buffers on one of the ports with Enterprise license	4095	
Total number of BB_credit buffers per module	5488	
Default Performance buffers	145	12
Total number of performance buffers per module	512 (shared)	

1. ISL = E port or TE port.
2. Ports on the 4-port 10-Gbps cannot operate in FL port mode.

Note The ports in the 4-port 10-Gbps switching module only support 10-Gbps dedicated rate mode. FL port mode and shared rate mode are not supported.

The following considerations apply to BB_credit buffers on 4-port 10-Gbps switching modules:

- BB_credit buffers for ISL connections can be configured from a minimum of 2 buffers to a maximum of 750 buffers.
- BB_credit buffers for Fx port mode connections can be configured from a minimum of 2 buffers to a maximum of 750 buffers.
- By default, 512 performance buffers are preallocated and are shared by all the ports. These buffers are configurable and the buffers are assigned to the port based on the availability of the buffers in the shared pool.
- There are 2488 extra buffers available as extended BB_credits after allocating all the default BB_credit buffers for all the ports in ISL mode (5488 - (750 * 4)).

Note Extended BB_credits are allocated across all ports on the switch. That is, they are not allocated by port group.

BB_Credit Buffers for Fabric Switches

This section describes how buffer credits are allocated to Cisco MDS 9000 Fabric switches, and includes the following topics:

- Cisco MDS 9148 Fabric Switch BB_Credit Buffers, page 4-16
- Cisco MDS 9148 Fabric Switch BB_Credit Buffers, page 4-16
- Cisco MDS 9124 Fabric Switch BB_Credit Buffers, page 4-17
- Cisco MDS 9222i Multiservice Modular Switch BB_Credit Buffers, page 4-17

Cisco MDS 9148 Fabric Switch BB_Credit Buffers

Table 4-10 lists the BB_credit buffer allocation for 48-port 8-Gbps Fibre Channel switches.

Table 4-10 48-Port 8-Gbps Fabric Switch BB_Credit Buffer Allocation

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port	
		ISL[1]	Fx Port
Default BB_credit buffers	128	32	32
Maximum configurable BB_credit buffers on 8-Gbps mode	128	125	125

1. ISL = E port or TE port.

The following considerations apply to BB_credit buffers on 48-port 8-Gbps Fabric Switches:

- BB_credit buffers can be configured from a minimum of 1 buffer to a maximum of 32 buffers per port when the ports are in F or FL mode.
- BB_credit buffers can be configured from a minimum of 2 buffers to a maximum of 32 buffers per port when the ports are in E or TE mode.

Cisco MDS 9134 Fabric Switch BB_Credit Buffers

Table 4-11 lists the BB_credit buffer allocation for MDS 9134 Fabric Switches.

Table 4-11 MDS 9134 Fabric Switch BB_Credit Buffer Allocation

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port	
		ISL[1]	Fx Port
Maximum user-configurable BB_credit buffers	64	61	61
Minimum user-configurable BB_credit buffers		2	1
Default BB_credit buffers on 10-Gbps mode	64	64	64
Default BB_credit buffers on 4-Gbps mode	64	16	16

1. ISL = E port or TE port.

Cisco MDS 9124 Fabric Switch BB_Credit Buffers

Table 4-12 lists the BB_credit buffer allocation for MDS 9124 Fabric Switches.

Table 4-12 MDS 9124 Fabric Switch BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL[1]	Fx Port
Maximum user-configurable BB_credit buffers	64	61	61
Minimum user-configurable BB_credit buffers		2	1
Default BB_credit buffers	64	16	16

1. ISL = E port or TE port.

Cisco MDS 9222i Multiservice Modular Switch BB_Credit Buffers

Table 4-13 lists the BB_credit buffer allocation for 18-port 4-Gbps Multiservice Modular switches.

Table 4-13 18-Port 4-Gbps Fabric Switch BB_Credit Buffer Allocation Defaults

BB_Credit Buffer Allocation Type	BB_Credit Buffers Per Port Group	BB_Credit Buffers Per Port Defaults	
		ISL[1]	Fx Port
User-configurable BB_credit buffers	4509	250	16

1. ISL = E port or TE port.

Extended BB_Credits

To facilitate BB_credits for long-haul links, the extended BB_credits feature allows you to configure the receive buffers above the maximum value on all Generation 2, Generation 3 and Generation 4 switching modules. When necessary, you can reduce the buffers on one port and assign them to another port, exceeding the default maximum. The minimum extended BB_credits per port is 256 and the maximum is 4095.

Note Extended BB_credits are not supported on the Cisco MDS 9148 Fabric Switch, Cisco MDS 9134 Fabric Switch, Cisco MDS 9124 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

In general, you can configure any port in a port group to dedicated rate mode. To do this, you must first release the buffers from the other ports before configuring larger extended BB_credits for a port.

Note The ENTERPRISE_PKG license is required to use extended BB_credits on Generation 2, Generation 3 and Generation 4 switching modules. Also, extended BB_credits are not supported by ports in shared rate mode.

All ports on the Generation 2 and Generation 3 switching modules support extended BB_credits. There are no limitations for how many extended BB_credits you can assign to a port (except for the maximum and minimum limits). If necessary, you can take interfaces out of service to make more extended BB_credits available to other ports.

You can use the extended BB_credits flow control mechanism in addition to BB_credits for long-haul links.

This section includes the following topics:

- Extended BB_credits on Generation 1 Switching Modules, page 4-18
- Extended BB_credits on Generation 2 and Generation 3 Switching Modules, page 4-19

Extended BB_credits on Generation 1 Switching Modules

The BB_credits feature allows you to configure up to 255 receive buffers on Generation 1 switching modules. To facilitate BB_credits for long haul links, you can configure up to 3,500 receive BB_credits on a Fibre Channel port on a Generation 1 switching module.

To use this feature on Generation 1 switching modules, you must meet the following requirements:

- Obtain the ENTERPRISE_PKG license. See the *Cisco MDS 9000 Family NX-OS Licensing Guide*.
- Configure this feature in any port of the full-rate 4-port group in either the Cisco MDS 9216i Switch or in the MPS-14/2 module (see Figure 4-9).

Figure 4-9 Port Group Support for the Extended BB_Credits Feature

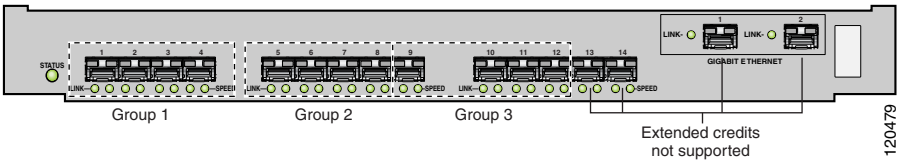

The port groups that support extended credit configurations are as follows:

- Any one port in ports 1 to 4 (identified as Group 1).
- Any one port in ports 5 to 8 (identified as Group 2).
- Any one port in ports 9 to 12 (identified as Group 3).

Note The last two Fibre Channel ports (port 13 and port 14) and the two Gigabit Ethernet ports do not support the extended BB_credits feature.

- Explicitly enable this feature in the required Cisco MDS switch.
- Disable the remaining three ports in the 4-port group if you need to assign more than 2,400 BB_credits to the first port in the port group.
  - If you assign less than 2,400 extended BB_credits to any one port in a port group, the remaining three ports in that port group can retain up to 255 BB_credits based on the port mode.

Note The receive BB_credit value for the remaining three ports depends on the port mode. The default value is 16 for the Fx mode and 255 for E or TE modes. The maximum value is 255 in all modes. This value can be changed as required without exceeding the maximum value of 255 BB_credits.

  - If you assign more than 2,400 (up to a maximum of 3,500) extended BB_credits to the port in a port group, you must disable the other three ports.

- If you change the BB_credit value the port is disabled, and then reenabled.
  - Disable (explicitly) this feature if you need to nondisruptively downgrade to Cisco SAN-OS Release 1.3 or earlier. When you disable this feature, the existing extended BB_credit configuration is completely erased.

Note The extended BB_credit configuration takes precedence over the receive BB_credit and performance buffer configurations.

Extended BB_credits on Generation 2 and Generation 3 Switching Modules

To use this feature on Generation 2 or Generation 3 switching modules, you must meet the following requirements:

- Display the interface configuration in the Information pane.
- Obtain the Enterprise package (ENTERPRISE_PKG) license (see the *NX-OS Family Licensing Guide*).
- Configure this feature in any port on a Generation 2 switch module. See the "Extended BB_Credits" section on page 4-17 for more information on extended BB_credits on Generation 2 switching modules.

Note Extended BB_credits are not supported on the Cisco MDS 9124 Fabric Switch, Cisco MDS 9134 Fabric Switch, the Cisco Fabric Switch for HP c-Class BladeSystem, and the Cisco Fabric Switch for IBM BladeCenter.

Buffer-to-Buffer Credit Recovery

Although the Fibre Channel standards require low bit error rates, bit errors do occur. Over time, the corruption of receiver-ready messages, known as R_RDY primitives, can lead to a loss of credits, which can eventually cause a link to stop transmitting in one direction. The Fibre Channel standards provide a feature for two attached ports to detect and correct this situation. This feature is called buffer-to-buffer credit recovery.

Buffer-to-buffer credit recovery functions as follows: the sender and the receiver agree to send checkpoint primitives to each other, starting from the time that the link comes up. The sender sends a checkpoint every time it has sent the specified number of frames, and the receiver sends a checkpoint every time it has sent the specified number of R_RDY primitives. If the receiver detects lost credits, it can retransmit them and restore the credit count on the sender.

The buffer-to-buffer credit recovery feature can be used on any non arbitrated loop link. This feature is most useful on unreliable links, such as MANs or WANs, but can also help on shorter, high-loss links, such as a link with a faulty fiber connection.

Note The buffer-to-buffer credit recovery feature is not compatible with the distance extension (DE) feature, also known as buffer-to-buffer credit spoofing. If you use intermediate optical equipment, such as DWDM transceivers or Fibre Channel bridges, on ISLs between switches that use DE, then buffer-to-buffer credit recovery on both sides of the ISL needs to be disabled.

Buffer-to-Buffer State Change Number

The BB_SC_N field (word 1, bits 15-12) specifies the buffer-to-buffer state change (BB_SC) number. The BB_SC_N field indicates that the sender of the port login (PLOGI), fabric login (FLOGI), or ISLs (E or TE ports) frame is requesting 2^SC_BB_N number of frames to be sent between two consecutive BB_SC send primitives, and twice the number of R_RDY primitives to be sent between two consecutive BB_SC receive primitives.

For Generation 2 and Generation 3 modules, the BB_SCN on ISLs (E or TE ports) is enabled by default. This can fail the ISLs if used with optical equipment using distance extension (DE), also known as buffer-to-buffer credit spoofing.

On a Generation-2 module, one port will not come up for the following configuration for all ports:

- Port Mode: auto or E for all the ports
- Rate Mode: dedicated
- Buffer Credits: default value

On a Generation-3 module, one or two ports will not come up for the following configuration for the first half of the ports, the second half of the ports, or all ports:

- Port Mode: auto or E for the first half of the ports, the second half of the ports, or for all of the ports
- Rate Mode: dedicated
- Buffer Credits: default value

When you configure port mode to auto or E and rate-mode to dedicated for all ports in the global buffer pool, you need to reconfigure buffer credits on one or more ports (other than default).

Note If you use distance extension (buffer-to-buffer credit spoofing) on ISLs between switches, the BB_SCN parameter on both sides of the ISL needs to be disabled.

Receive Data Field Size

You can also configure the receive data field size for Fibre Channel interfaces. If the default data field size is 2112 bytes, the frame length will be 2148 bytes.

Configuring Interface Buffers

This section includes the following topics:

- Configuring Buffer-to-Buffer Credits, page 4-21
- Configuring Performance Buffers, page 4-22
- Configuring Extended BB_credits, page 4-23
- Enabling Buffer-to-Buffer Credit Recovery, page 4-24
- Enabling the Buffer-to-Buffer State Change Number, page 4-24
- Configuring Receive Data Field Size, page 4-24

Configuring Buffer-to-Buffer Credits

Detailed Steps

To configure BB_credits for a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.

	Command	Purpose
Step 3	`switch(config-if)# switchport fcrxbbcredit default`	Applies the default operational value to the selected interface. The operational value depends on the port mode. The default values are assigned based on the port capabilities.
	`switch(config-if)# switchport fcrxbbcredit 5`	Assigns a BB_credit of 5 to the selected interface. The range to assign BB_credits is between 1 and 255.
	`switch(config-if)# switchport fcrxbbcredit 5 mode E`	Assigns this value if the port is operating in E or TE mode. The range to assign BB_credits is between 1 and 255.
	`switch(config-if)# switchport fcrxbbcredit 5 mode Fx`	Assigns this value if the port is operating in F or FL mode. The range to assign BB_credits is between 1 and 255.
Step 4	`switch(config-if# do show int fc1/1`	Displays the receive and transmit BB_credit along with other pertinent interface information for this interface. **Note** The BB_credit values are correct at the time the registers are read. They are useful to verify situations when the data traffic is slow.

This example shows the output of the **show int fc1/1** command:

```
intfc1/1 is up
    ...
      16 receive B2B credit remaining
      3 transmit B2B credit remaining
```

Configuring Performance Buffers

Detailed Steps

To configure performance buffers for a Fibre Channel interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcrxbbcredit performance-buffers 45`	Assigns a performance buffer of 45 to the selected interface. The value ranges from 1 to 145.
	`switch(config-if)# switchport fcrxbbcredit performance-buffers default`	Reverts to the factory default of using the built-in algorithm.

Note Use the **show interface bbcredit** command to display performance buffer values and other BB_credit information.

Configuring Extended BB_credits

Detailed Steps

To configure extended BB_credits for a MDS-14/2 interface, for a Generation 2 switching module interface (not including the Cisco MDS 9124 Fabric Switch), or for an interface in a Cisco MDS 9216i switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# fcrxbbcredit extended enable`	Enables the extended BB_credits feature.
	`switch(config)# no fcrxbbcredit extended enable`	Disables (default) the extended BB_credits feature.
Step 3	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 4	`switch(config-if)# switchport fcrxbbcredit extended 1500`	Applies the extended BB_credit value of 1,500 credits to the selected interface. The valid range is from 256 to 3,500 credits.
	`switch(config-if)# no switchport fcrxbbcredit extended 1500`	Clears the configured extended BB_credit configuration for this port.

Step 5	`switch# do show interface fc3/2`
	`fc3/2 is trunking`
	`Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)`
	`Port WWN is 20:82:00:05:30:00:2a:1e`
	`Peer port WWN is 20:42:00:0b:46:79:f1:80`
	`Admin port mode is auto, trunk mode is on`
	`Port mode is TE`
	`Port vsan is 1`
	`Speed is 2 Gbps`
	`Transmit B2B Credit is 255`
	`Receive B2B Credit is 1500`
	`Receive data field Size is 2112`
	`...`
		Displays the receive and transmit BB_credit values along with other pertinent interface information for this interface if the interface is in the up state.
		Note The receive BB_credit value reflects the extended BB_credit configuration, if applicable.

Enabling Buffer-to-Buffer Credit Recovery

Buffer-to-buffer credit recovery on ISLs (E or TE ports) is enabled by default.

Detailed Steps

To use buffer-to-buffer credit recovery on a port, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcbbscn`	Enables buffer-to-buffer credit recovery on the interface.
	`switch(config-if)# no switchport fcbbscn`	Disables (default) buffer-to-buffer credit recovery on the interface.

Enabling the Buffer-to-Buffer State Change Number

Detailed Steps

To use the BB_SC_N field during PLOGI or FLOGI, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc 1/1` `switch(config-if)#`	Selects the interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcbbscn`	Enables the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.
	`switch(config-if)# no switchport fcbbscn`	Disables (default) the use of buffer-to-buffer state change number for PLOGIs and FLOGIs on the interface.

Configuring Receive Data Field Size

Detailed Steps

To configure the receive data field size, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Selects a Fibre Channel interface and enters interface configuration submode.
Step 3	`switch(config-if)# switchport fcrxbufsize 2000`	Reduces the data field size for the selected interface to 2000 bytes. The default is 2112 bytes and the range is from 256 to 2112 bytes.

Verifying BB_Credit Configuration

To display BB_credit configuration information, perform one of the following tasks:

Command	Purpose
show interface	Displays the interface configuration.
show interface bbcredit	Displays the BB_credit configuration for all the interfaces.
show interface fc *numbers* **bbcredit**	Displays the BB_credit configuration for the specified interfaces.

For detailed information about the fields in the output from these commands, refer to the *Cisco NX-OS Command Reference*.

To display the BB_credit information, use the **show interface bbcredit** command (see Example 4-1 and Example 4-2).

Example 4-1 Displays BB_credit Information

```
switch# show interface bbcredit
fc2/1 is down (SFP not present)
...
fc2/17 is trunking
    Transmit B2B Credit is 255
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 375
      12 receive B2B credit remaining
      255 transmit B2B credit remaining
fc2/18 is down (SFP not present)
fc2/19 is down (SFP not present)
fc2/20 is down (SFP not present)
fc2/21 is down (Link failure or not-connected)
...
fc2/31 is up
    Transmit B2B Credit is 0
    Receive B2B Credit is 12
    Receive B2B Credit performance buffers is 48
      12 receive B2B credit remaining
      0 transmit B2B credit remaining
fc2/32 is down (Link failure or not-connected)
```

Example 4-2 Displays BB_credit Information for a Specified Fibre Channel Interface

```
switch# show interface fc2/31 bbcredit
fc2/31 is up
```

```
Transmit B2B Credit is 0
Receive B2B Credit is 12
Receive B2B Credit performance buffers is 48
  12 receive B2B credit remaining
  0 transmit B2B credit remaining
```

Configuring Trunking

- Information About Trunking, page 5-1
- Guidelines and Limitations, page 5-7
- Default Settings, page 5-11
- Configuring Trunking, page 5-11
- Verifying Trunking Configuration, page 5-13
- Configuration Example for F Port Trunking, page 5-14

Information About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports (See Figure 5-1 and Figure 5-2).

This section includes the following topics:

- Trunking E Ports, page 5-1
- Trunking F Ports, page 5-2
- Key Concepts, page 5-3
- Trunking Protocols, page 5-3
- Trunk Modes, page 5-4
- Trunk-Allowed VSAN Lists and VF_IDs, page 5-5

Trunking E Ports

Trunking the E ports enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

Figure 5-1 Trunking E Ports

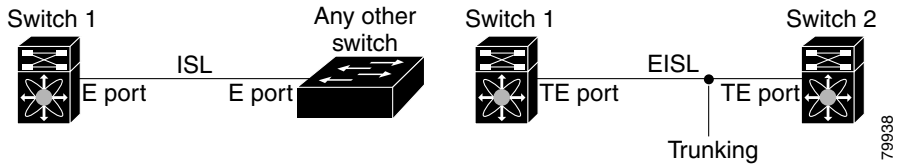

Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c_Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Trunking F Ports

Trunking F ports allows interconnected ports to transmit and receive tagged frames in more than one VSAN, over the same physical link.

Figure 5-2 represents the possible trunking scenarios in a SAN with MDS core switches, NPV switches, third-party core switches, and HBAs.

Figure 5-2 Trunking F Ports

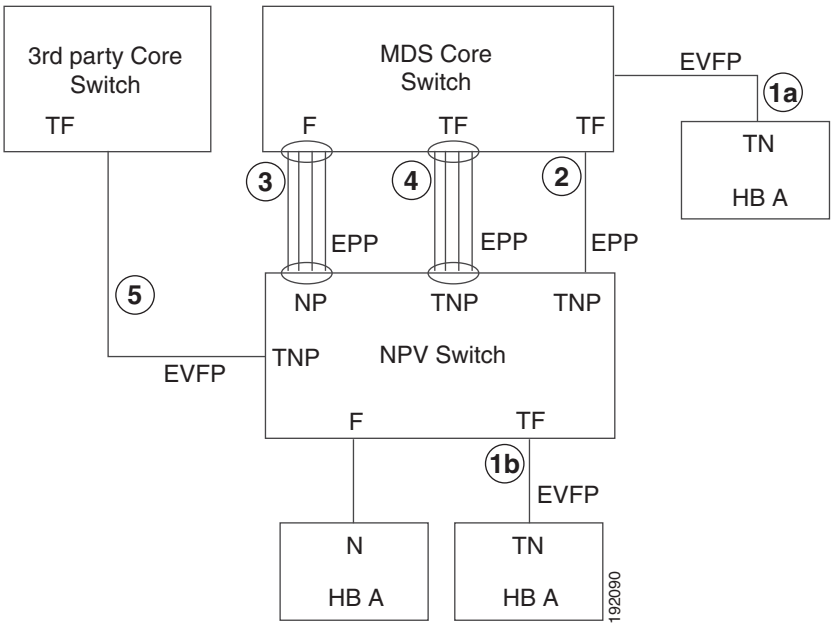

Link Number	Link Description
1a and 1b	F port trunk with N port.[1]
2	F port trunk with NP port.

Link Number	Link Description
3	F PortChannnel with NP port.
4	Trunked F PortChannel with NP port.
5	Trunking NP port with third-party core switch F port.[1]

1. These features are not supported currently.

Key Concepts

The trunking feature includes the following key concepts:

- TE port—If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- TF port—If trunk mode is enabled in an F port (see the link 2 in Figure 5-2) and that port becomes operational as a trunking F port, it is referred to as a TF port.
- TN port—If trunk mode is enabled (not currently supported) in an N port (see the link 1b in Figure 5-2) and that port becomes operational as a trunking N port, it is referred to as a TN port.
- TNP port—If trunk mode is enabled in an NP port (see the link 2 in Figure 5-2) and that port becomes operational as a trunking NP port, it is referred to as a TNP port.
- TF PortChannel—If trunk mode is enabled in an F PortChannel (see the link 4 in Figure 5-2) and that PortChannel becomes operational as a trunking F PortChannel, it is referred to as TF PortChannel. Cisco Port Trunking Protocol (PTP) is used to carry tagged frames.
- TF-TN port link—A single link can be established to connect an F port to an HBA to carry tagged frames (see the link 1a and 1b in Figure 5-2) using Exchange Virtual Fabrics Protocol (EVFP). A server can reach multiple VSANs through a TF port without inter-VSAN routing (IVR).
- TF-TNP port link—A single link can be established to connect an TF port to an TNP port using the PTP protocol to carry tagged frames (see the link 2 in Figure 5-2). PTP is used because PTP also supports trunking PortChannels.

  Note The TF-TNP port link between a third-party NPV core and a Cisco NPV switch is established using the EVFP protocol.

- A Fibre Channel VSAN is called Virtual Fabric and uses a VF_ID in place of the VSAN ID. By default, the VF_ID is 1 for all ports. When an N port supports trunking, a pWWN is defined for each VSAN and called a logical pWWN. In the case of MDS core switches, the pWWNs for which the N port requests additional FC_IDs are called virtual pWWNs.

Trunking Protocols

The trunking protocol is important for trunking operations on the ports. The protocols enable the following activities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

Table 5-1 specifies the protocols used for trunking and channeling.

Table 5-1 Supported Trunking Protocols

Trunk Link	Default
TE-TE port link	Cisco EPP (PTP)
TF-TN port link[1]	FC-LS Rev 1.62 EVFP
TF-TNP port link	Cisco EPP (PTP)
E or F PortChannel	Cisco EPP (PCP)
TF Port Channel	Cisco EPP (PTP and PCP)
Third-party TF-TNP port link[1]	FC-LS Rev 1.62 EVFP

1. These features are not currently supported.

By default, the trunking protocol is enabled on E ports and disabled on F ports. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected. The TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

Note We recommend that both ends of a trunking link belong to the same port VSAN. On certain switches or fabric switches where the port VSANs are different, one end returns an error and the other end is not connected.

Trunk Modes

By default, trunk mode is enabled on all Fibre Channel interfaces (Mode: E, F, FL, Fx, ST, and SD) on non-NPV switches. On NPV switches, by default, trunk mode is disabled. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The trunk mode configuration at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends (see Table 5-2).

Table 5-2 Trunk Mode Status Between Switches

Your Trunk Mode Configuration			Resulting State and Port Mode	
Port Type	**Switch 1**	**Switch 2**	**Trunking State**	**Port Mode**
E ports	On	Auto or on	Trunking (EISL)	TE port
	Off	Auto, on, or off	No trunking (ISL)	E port
	Auto	Auto	No trunking (ISL)	E port
Port Type	**Core Switch**	**NPV Switch**	**Trunking State**	**Link Mode**
F and NP ports	On	Auto or on	Trunking	TF-TNP link
	Auto	On	Trunking	TF-TNP link
	Off	Auto, on, or off	No trunking	F-NP link

Tip The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other side set to on.

Note When connected to a third-party switch, the trunk mode configuration on E ports has no effect. The ISL is always in a trunking disabled state. In the case of F ports, if the third-party core switch ACC's physical FLOGI with the EVFP bit is configured, then EVFP protocol enables trunking on the link.

Trunk-Allowed VSAN Lists and VF_IDs

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

Switch 1 (see Figure 5-3) has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational (see Figure 5-3).

For all F, N, and NP ports, the default VF_ID is 1 when there is no VF_ID configured. The trunk-allowed VF_ID list on a port is same as the list of trunk-allowed VSANs. VF_ID 4094 is called the control VF_ID and it is used to define the list of trunk-allowed VF-IDs when trunking is enabled on the link.

If F port trunking and channeling is enabled, or if **switchport trunk mode on** is configured in NPV mode for any interface, or if NP PortChannel is configured, the VSAN and VF-ID ranges available for the configuration are as described in Table 5-3.

Table 5-3 VSAN and VF-ID Reservations

VSAN or VF-ID	Description
000h	Cannot be used as virtual fabric identifier.
001h(1) to EFFh(3839)	This VSAN range is available for user configuration.
F00h(3840) to FEEh(4078)	Reserved VSANs and they are not available for user configuration.
FEFh(4079)	EVFP isolated VSAN.
FF0h(4080) to FFEh(4094)	Used for vendor-specific VSANs.
FFFh	Cannot be used as virtual fabric identifier.

Note If the VF_ID of the F port and the N port do not match, then no tagged frames can be exchanged.

Figure 5-3 Default Allowed-Active VSAN Configuration

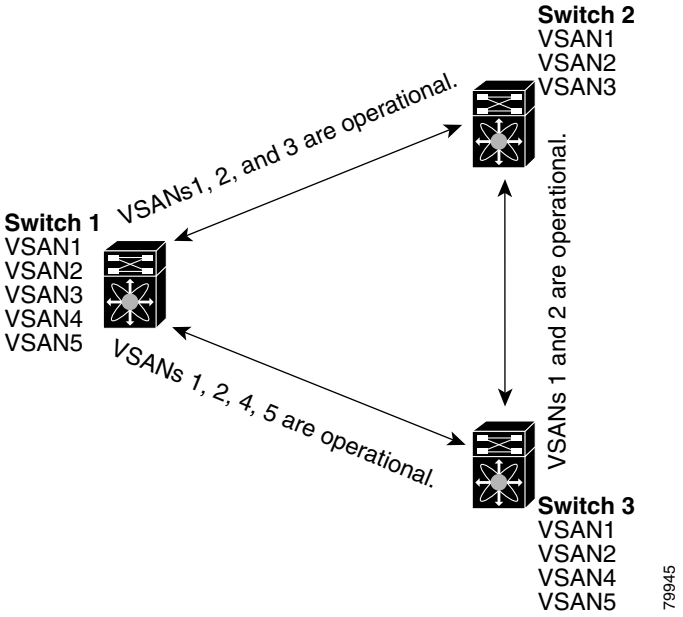

You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using Figure 5-3 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 5-4). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Figure 5-4 Operational and Allowed VSAN Configuration

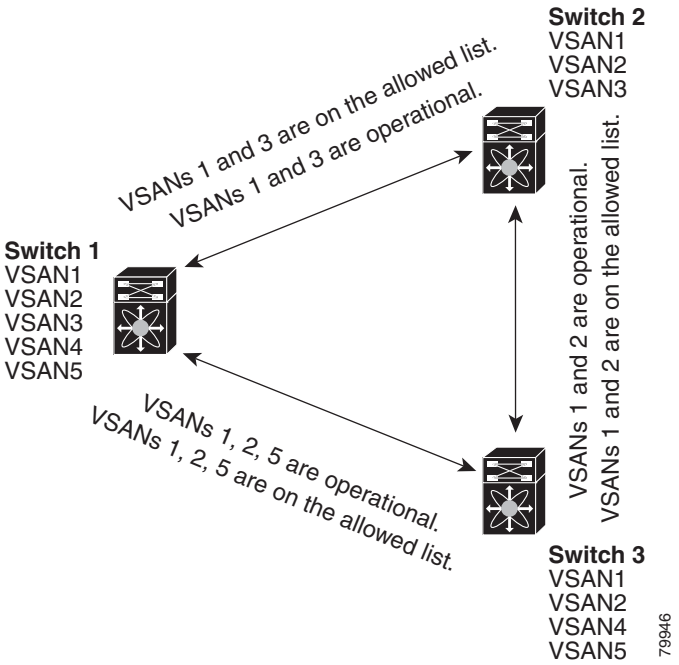

Guidelines and Limitations

Trunking has the following configuration guidelines and limitations:

- General Guidelines and Limitations, page 5-7
- Upgrade and Downgrade Limitations, page 5-8
- Difference Between TE Ports and TF-TNP Ports, page 5-8
- Trunking Misconfiguration Examples, page 5-10

General Guidelines and Limitations

The trunking feature has the following general configuration guidelines and limitations:

- F ports support trunking in Fx mode.
- The trunk-allowed VSANs configured for TE, TF, and TNP links are used by the trunking protocol to determine the allowed active VSANs in which frames can be received or transmitted.
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.
- Trunking F ports and trunking F PortChannels are not supported on the following hardware:
  - 91x4 switches, if NPIV is enabled and used as the NPIV core switch.
  - Generation 1 2-Gbps Fibre Channel switching modules.
- On core switches, the FC-SP authentication will be supported only for the physical FLOGI from the physical pWWN.

- No FC-SP authentication is supported by the NPV switch on the server F ports.
- MDS does not enforce the uniqueness of logical pWWNs across VSANs.
- DPVM is not supported on trunked F port logins.
- The DPVM feature is limited to the control of the port VSAN, since the EVFP protocol does not allow changing the VSAN on which a logical pWWN has done FLOGI.
- The port security configuration will be applied to both the first physical FLOGI and the per VSAN FLOGIs.
- Trunking is not supported on F ports that have FlexAttach enabled.
- On MDS 91x4 core switches, hard zoning can be done only on F ports that are doing either NPIV or trunking. However, in NPV mode, this restriction does not apply since zoning is enforced on the core F port.

Note Fibre Channel Security Protocol (FC-SP) is not supported for 6.2(1) release on MDS 9710, but targeted for a future release.

Upgrade and Downgrade Limitations

The trunking and channeling feature includes the following upgrade and downgrade limitations:

- When F port trunking or channeling is configured on a link, the switch cannot be downgraded to Cisco MDS SAN-OS Release 3.x and NX-OS Release 4.1(1b), or earlier.
- If you are upgrading from a SAN-OS Release 3.x to NX-OS Release 5.0(1), and you have not created VSAN 4079, the NX-OS software will automatically create VSAN 4079 and reserve it for EVFP use.

  If VSAN 4079 is reserved for EVFP use, the **switchport trunk allowed vsan** command will filter out VSAN 4079 from the allowed list, as shown in the following example:

```
switch(config-if)# switchport trunk allowed vsan 1-4080
1-4078,4080
switch(config-if)#
```

  - If you have created VSAN 4079, the upgrade to NX-OS Release 5.0(1) will have no affect on VSAN 4079.
  - If you downgrade after NX-OS Release 5.0(1), the VSAN will no longer be reserved for EVFP use.

Difference Between TE Ports and TF-TNP Ports

In case of TE ports, the VSAN will in be initializing state when VSAN is coming up on that interface and when peers are in negotiating phase. Once the handshake is done, VSAN will be moved to up state in the successful case, and isolated state in the case of failure. Device Manager will show the port status as amber during initializing state and it will be green once VSANs are up.

This example shows the trunk VSAN states of a TE port:

```
Switch# show interface fc2/15
fc2/15 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
```

```
    Port WWN is 20:4f:00:0d:ec:6d:2b:40
    Peer port WWN is 20:0a:00:0d:ec:3f:ab:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 250
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (isolated)                 (100-101)
    Trunk vsans (initializing)             ()
```

In case of TF ports, after the handshake, one of the allowed VSANs will be moved to the up state. All other VSANs will be in initializing state even though the handshake with the peer is completed and successful. Each VSAN will be moved from initializing state to up state when a server or target logs in through the trunked F or NP ports in the corresponding VSAN.

Note In case of TF or TNP ports, the Device Manager will show the port status as amber even after port is up and there is no failure. It will be changed to green once all the VSAN has successful logins.

This example shows a TF port information after the port is in the up state:

```
sw7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

This example shows the TF port information when a server logs in on noninternal FLOGI VSAN. VSAN 2183 is moved to the up state when the server logs in to VSAN 2183.

```
w7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
```

```
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2
182-2183)
    Trunk vsans (up)                       (1,2183)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

Trunking Misconfiguration Examples

If you do not configure the VSANs correctly, issues with the connection may occur. For example, if you merge the traffic in two VSANs, both VSANs will be mismatched. The trunking protocol validates the VSAN interfaces at both ends of a link to avoid merging VSANs (see Figure 5-5).

Figure 5-5 VSAN Mismatch

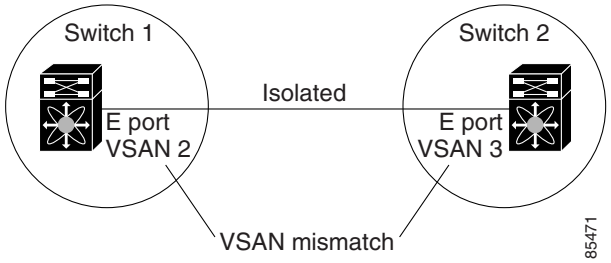

The trunking protocol detects potential VSAN merging and isolates the ports involved (see Figure 5-5).

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see Figure 5-6).

Figure 5-6 Third-Party Switch VSAN Mismatch

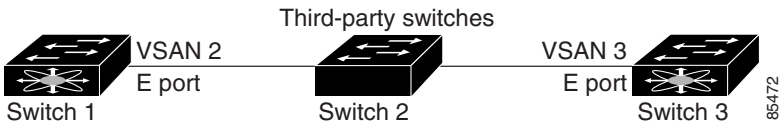

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. Cisco DCNM-SAN helps detect such topologies.

Default Settings

Table 5-4 lists the default settings for trunking parameters.

Table 5-4 Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	ON on non-NPV and MDS core switches. OFF on NPV switches.
Allowed VSAN list	1 to 4093 user-defined VSAN IDs.
Allowed VF-ID list	1 to 4093 user-defined VF-IDs.
Trunking protocol on E ports	Enabled.
Trunking protocol on F ports	Disabled.

Configuring Trunking

This section includes the following topics:

- Enabling the Cisco Trunking and Channeling Protocols, page 5-11
- Enabling the F Port Trunking and Channeling Protocol, page 5-12
- Configuring Trunk Mode, page 5-12
- Configuring an Allowed-Active List of VSANs, page 5-12

Enabling the Cisco Trunking and Channeling Protocols

This section describes how to enable the required trunking and channeling protocols.

Prerequisites

- To avoid inconsistent configurations, disable all ports with a **shutdown** command before enabling or disabling the trunking protocols.

Detailed Steps

To enable or disable the Cisco trunking and channeling protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# trunk protocol enable` `switch(config)#`	Enables the Cisco PTP trunking protocol (default).
	`switch(config)# no trunk protocol enable` `switch(config)#`	Disables the Cisco PTP trunking protocol.

Enabling the F Port Trunking and Channeling Protocol

This section describes how to enable the F port trunking and channeling protocol.

Prerequisites

- To avoid inconsistent configurations, shut all ports before enabling or disabling the trunking protocols.

Detailed Steps

To enable or disable the F port trunking and channeling protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config tasf`	Enters configuration mode.
Step 2	`switch(config)# feature fport-channel-trunk` `switch(config)#`	Enables the F port trunking and channeling protocol (default).
	`switch(config)# no feature fport-channel-trunk` `switch(config)#`	Disables the F port trunking and channeling protocol.

Configuring Trunk Mode

Detailed Steps

To configure trunk mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# switchport trunk mode on`	Enables (default) the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode off`	Disables the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode auto`	Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

Configuring an Allowed-Active List of VSANs

Detailed Steps

To configure an allowed-active list of VSANs for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.

	Command	Purpose
Step 3	`switch(config-if)# switchport trunk allowed vsan 2-4`	Changes the allowed list for the specified VSANs.
	`switch(config-if)# switchport trunk allowed vsan add 5` `updated trunking membership`	Expands the specified VSAN (5) to the new allowed list.
	`switch(config-if)# no switchport trunk allowed vsan 2-4`	Deletes VSANs 2, 3, and 4.
	`switch(config-if)# no switchport trunk allowed vsan add 5`	Deletes the expanded allowed list.

Verifying Trunking Configuration

To display trunking configuration information, perform one of the following tasks:

Command	Purpose
show interface fc *slot/port*	Displays the interface configuration information that includes trunking, trunk mode, allowed VSANs, and status.
show trunk protocol	Displays whether the trunk protocol is enabled.
show interface trunk vsan *numbers*	Displays whether the interface is trunking, and the allowed VSAN list for each trunking interface.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples 5-1 to 5-3.

Example 5-1 Displays a Trunked Fibre Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
    Hardware is Fibre Channel
    Port WWN is 20:0d:00:05:30:00:58:1e
    Peer port WWN is 20:0d:00:05:30:00:59:1e
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Receive B2B Credit is 255
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      233996 frames input, 14154208 bytes, 0 discards
        0 CRC, 0 unknown class
        0 too long, 0 too short
      236 frames output, 13818044 bytes, 0 discards
      11 input OLS, 12 LRR, 10 NOS, 28 loop inits
```

```
      34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

Example 5-2 Displays the Trunking Protocol

```
switch# show trunk protocol
Trunk protocol is enabled
```

Example 5-3 Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
    Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
    Vsan 1 is up, FCID is 0x760001
    Vsan 2 is up, FCID is 0x6f0001

fc3/11 is trunking
    Belongs to port-channel 6
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
```

Configuration Example for F Port Trunking

This example shows how to configure trunking and bring up the TF-TNP link between an F port in the NPIV core switch and an NP port in the NPV switch:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

```
switch(config)# feature npiv
```

Step 3 Configure the port mode to auto, F, or Fx on the MDS core switch:

```
switch(config)# interface fc1/2
switch(config-if)# switchport mode F
```

Step 4 Configure the trunk mode to ON on the MDS core switch:

```
switch(config-if)# switchport trunk mode on
```

Step 5 Configure the port mode to NP on the NPV switch:

```
switch(config)# interface fc1/2
switch(config-if)# switchport mode NP
```

Step 6 Configure the trunk mode to ON on the NPV switch:

```
switch(config-if)# switchport trunk mode on
```

Step 7 Set the port administrative state on NPIV and NPV switches to ON:

```
switch(config)# interface fc1/2
switch(config-if)# shut
switch(config-if)# no shut
```

Step 8 Save the configuration.

```
switch(config)# copy running-config startup-config
```

Configuring PortChannels

- Information About PortChannels, page 6-1
- Prerequisites for PortChannels, page 6-12
- Guidelines and Limitations, page 6-13
- Default Settings, page 6-16
- Configuring PortChannels, page 6-17
- Verifying PortChannel Configuration, page 6-20
- Configuration Examples for F and TF PortChannels, page 6-24

Information About PortChannels

This section includes the following topics:

- PortChannels Overview, page 6-2
- E PortChannels, page 6-2
- F and TF PortChannels, page 6-3
- PortChanneling and Trunking, page 6-3
- Load Balancing, page 6-4
- PortChannel Modes, page 6-6
- PortChannel Deletion, page 6-7
- Interfaces in a PortChannel, page 6-7
- PortChannel Protocols, page 6-9
- Channel Group Creation, page 6-10
- Autocreation, page 6-11
- Manually Configured Channel Groups, page 6-12

PortChannels Overview

PortChannels refer to the aggregation of multiple physical interfaces into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy (See Figure 6-1). PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

Figure 6-1 PortChannel Flexibility

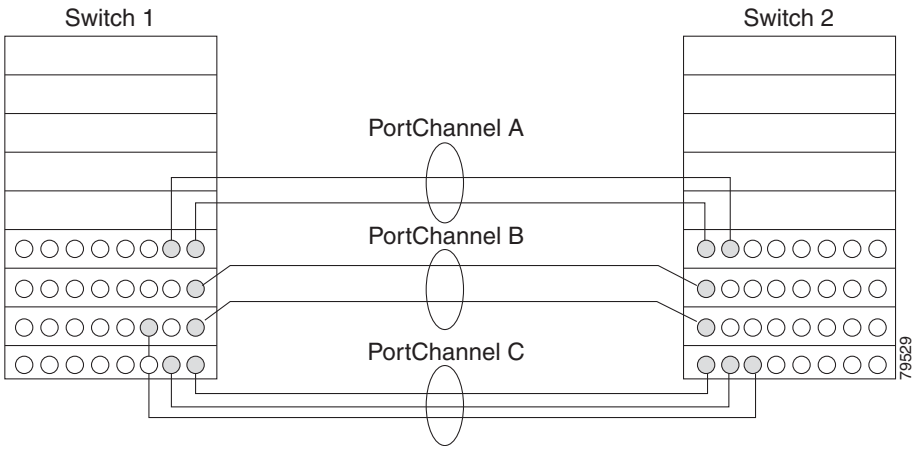

PortChannels on Cisco MDS 9000 Family switches allow flexibility in configuration. This illustrates three possible PortChannel configurations:

- PortChannel A aggregates two links on two interfaces on the same switching module at each end of a connection.
- PortChannel B also aggregates two links, but each link is connected to a different switching module. If the switching module goes down, traffic is not affected.
- PortChannel C aggregates three links. Two links are on the same switching module at each end, while one is connected to a different switching module on switch 2.

E PortChannels

An E PortChannel refers to the aggregation of multiple E ports into one logical interface to provide higher aggregated bandwidth, load balancing, and link redundancy. PortChannels can connect to interfaces across switching modules, so a failure of a switching module cannot bring down the PortChannel link.

A PortChannel has the following features and restrictions:

- Provides a point-to-point connection over ISL (E ports) or EISL (TE ports). Multiple links can be combined into a PortChannel.
- Increases the aggregate bandwidth on an ISL by distributing traffic among all functional links in the channel.
- Load balances across multiple links and maintains optimum bandwidth utilization. Load balancing is based on the source ID, destination ID, and exchange ID (OX ID).

- Provides high availability on an ISL. If one link fails, traffic previously carried on this link is switched to the remaining links. If a link goes down in a PortChannel, the upper protocol is not aware of it. To the upper protocol, the link is still there, although the bandwidth is diminished. The routing tables are not affected by link failure. PortChannels may contain up to 16 physical links and may span multiple modules for added high availability.

Note See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide* for information about failover scenarios for PortChannels and FSPF links.

F and TF PortChannels

An F PortChannel is also a logical interface that combines a set of F ports connected to the same Fibre Channel node and operates as one link between the F ports and the NP ports. The F PortChannels support bandwidth utilization and availability like the E PortChannels. F PortChannels are mainly used to connect MDS core and NPV switches to provide optimal bandwidth utilization and transparent failover between the uplinks of a VSAN.

An F PortChannel trunk combines the functionality and advantages of a TF port and an F PortChannel. This logical link uses the Cisco PTP and PCP protocols over Cisco EPP (ELS).

Note If a Cisco MDS 9124 or 9134 switch is used as a core switch, only a nontrunking F PortChannel is supported. Trunking is not supported on this platform when NPIV enabled.

PortChanneling and Trunking

Trunking is a commonly used storage industry term. However, the Cisco NX-OS software and switches in the Cisco MDS 9000 Family implement trunking and PortChanneling as follows:

- PortChanneling enables several physical links to be combined into one aggregated logical link.
- Trunking enables a link transmitting frames in the EISL format to carry (trunk) multiple VSAN traffic. For example, when trunking is operational on an E port, that E port becomes a TE port. A TE port is specific to switches in the Cisco MDS 9000 Family. An industry standard E port can link to other vendor switches and is referred to as a nontrunking interface (See Figure 6-2 and Figure 6-3). See Chapter 5, "Configuring Trunking," for information on trunked interfaces.

Figure 6-2 Trunking Only

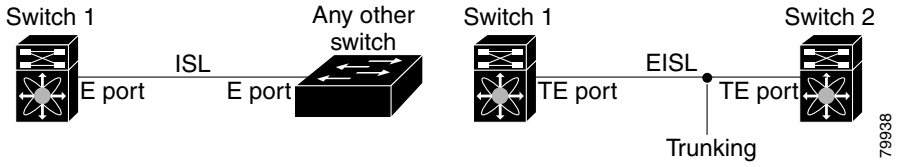

Figure 6-3 PortChanneling and Trunking

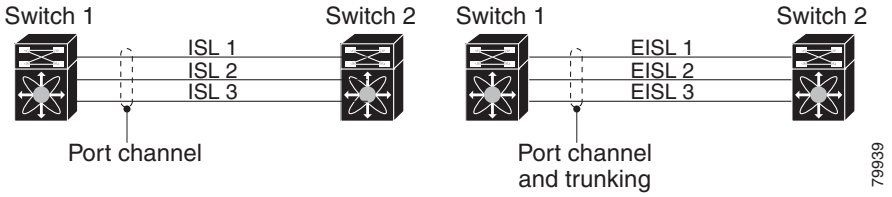

PortChanneling and trunking are used separately across an ISL.

- PortChanneling—Interfaces can be channeled between the following sets of ports:
  - E ports and TE ports
  - F ports and NP ports
  - TF ports and TNP ports
- Trunking—Trunking permits carrying traffic on multiple VSANs between switches.

  See the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.
- Both PortChanneling and trunking can be used between TE ports over EISLs.

Load Balancing

Two methods support the load-balancing functionality:

- Flow based—All frames between source and destination follow the same links for a given flow. That is, whichever link is selected for the first exchange of the flow is used for all subsequent exchanges.
- Exchange based—The first frame in an exchange picks a link and subsequent frames in the exchange follow the same link. However, subsequent exchanges can use a different link. This provides more granular load balancing while preserving the order of frames for each exchange.

Figure 6-4 illustrates how source ID 1 (SID1) and destination ID1 (DID1) based load balancing works. When the first frame in a flow is received on an interface for forwarding, link 1 is selected. Each subsequent frame in that flow is sent over the same link. No frame in SID1 and DID1 utilizes link 2.

Figure 6-4 SID1 and DID1-Based Load Balancing

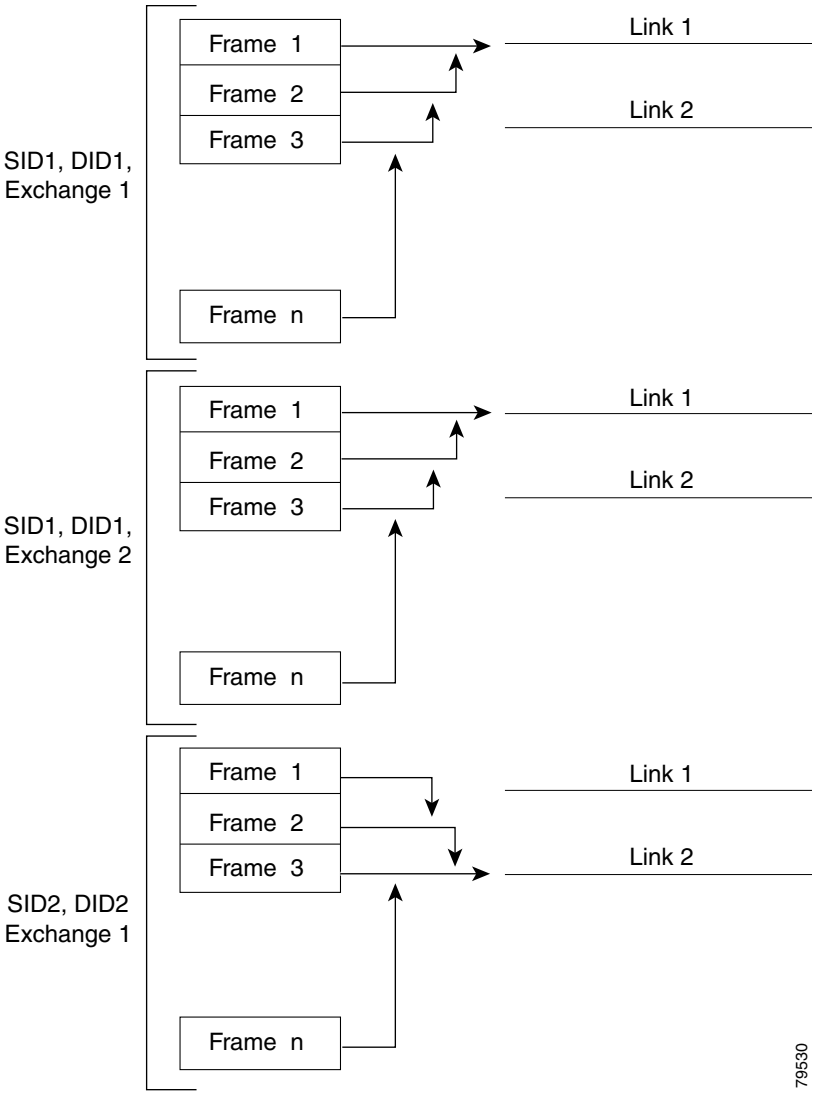

Figure 6-5 illustrates how exchange-based load balancing works. When the first frame in an exchange is received for forwarding on an interface, link 1 is chosen by a hash algorithm. All remaining frames in that particular exchange are sent on the same link. For exchange 1, no frame uses link 2. For the next exchange, link 2 is chosen by the hash algorithm. Now all frames in exchange 2 use link 2.

Figure 6-5 SID1, DID1, and Exchange-Based Load Balancing

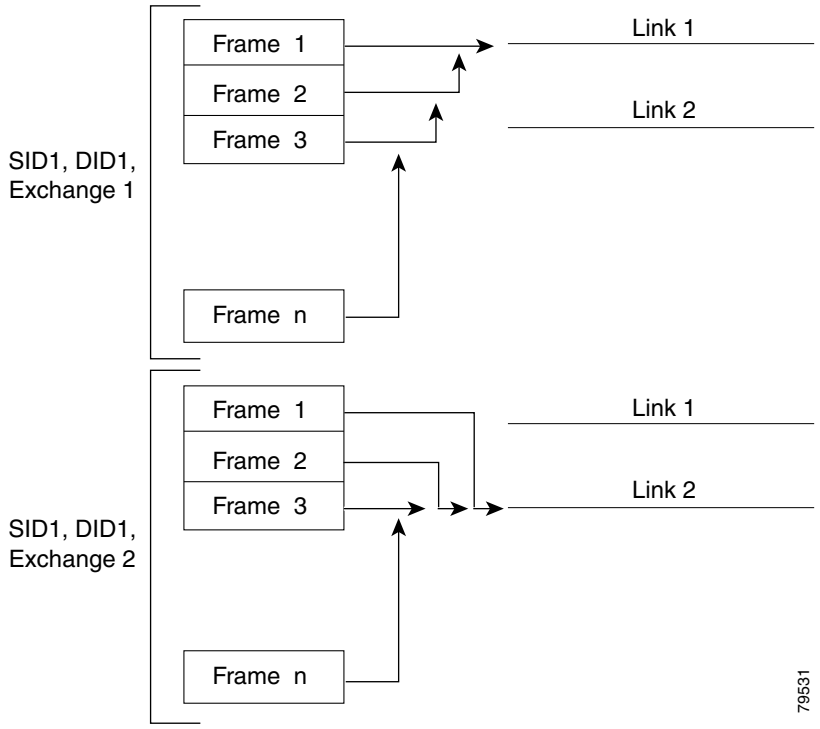

For more information on configuring load balancing and in-order delivery features, see the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

PortChannel Modes

You can configure each PortChannel with a channel group mode parameter to determine the PortChannel protocol behavior for all member ports in this channel group. The possible values for a channel group mode are as follows:

- ON (default)—The member ports only operate as part of a PortChannel or remain inactive. In this mode, the PortChannel protocol is not initiated. However, if a PortChannel protocol frame is received from a peer port, the software indicates its nonnegotiable status. This mode is backward compatible with the existing implementation of PortChannels in releases prior to Release 2.0(1b), where the channel group mode is implicitly assumed to be ON. In Cisco MDS SAN-OS Releases 1.3 and earlier, the only available PortChannel mode was the ON mode. PortChannels configured in the ON mode require you to explicitly enable and disable the PortChannel member ports at either end if you add or remove ports from the PortChannel configuration. You must physically verify that the local and remote ports are connected to each other.
- ACTIVE—The member ports initiate PortChannel protocol negotiation with the peer port(s) regardless of the channel group mode of the peer port. If the peer port, while configured in a channel group, does not support the PortChannel protocol, or responds with a nonnegotiable status, it will default to the ON mode behavior. The ACTIVE PortChannel mode allows automatic recovery without explicitly enabling and disabling the PortChannel member ports at either end.

Table 6-1 compares ON and ACTIVE modes.

Table 6-1 Channel Group Configuration Differences

ON Mode	ACTIVE Mode
No protocol is exchanged.	A PortChannel protocol negotiation is performed with the peer ports.
Moves interfaces to the suspended state if its operational values are incompatible with the PortChannel.	Moves interfaces to the isolated state if its operational values are incompatible with the PortChannel.
When you add or modify a PortChannel member port configuration, you must explicitly disable (shut) and enable (no shut) the PortChannel member ports at either end.	When you add or modify a PortChannel interface, the PortChannel automatically recovers.
Port initialization is not synchronized.	There is synchronized startup of all ports in a channel across peer switches.
All misconfigurations are not detected as no protocol is exchanged.	Consistently detect misconfigurations using a PortChannel protocol.
Transitions misconfigured ports to the suspended state. You must explicitly disable (shut) and enable (no shut) the member ports at either end.	Transitions misconfigured ports to the isolated state to correct the misconfiguration. Once you correct the misconfiguration, the protocol ensures automatic recovery.
This is the default mode.	You must explicitly configure this mode.

PortChannel Deletion

When you delete the PortChannel, the corresponding channel membership is also deleted. All interfaces in the deleted PortChannel convert to individual physical links. After the PortChannel is removed, regardless of the mode used (ACTIVE and ON), the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Graceful Shutdown" section on page 2-9).

If you delete the PortChannel for one port, then the individual ports within the deleted PortChannel retain the compatibility parameter settings (speed, mode, port VSAN, allowed VSAN, and port security). You can explicitly change those settings as required.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

Interfaces in a PortChannel

You can add or remove a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel. Removing an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

This section describes interface configuration for a PortChannel and includes the following topics:

- Interface Addition to a PortChannel, page 6-8

- Forcing an Interface Addition, page 6-9
- Interface Deletion from a PortChannel, page 6-9

Note For information about PortChannel support on Generation 2 switching modules, see the "PortChannel Limitations" section on page 3-21.

Interface Addition to a PortChannel

You can add a physical interface (or a range of interfaces) to an existing PortChannel. The compatible parameters on the configuration are mapped to the PortChannel. Adding an interface to a PortChannel increases the channel size and bandwidth of the PortChannel.

A port can be configured as a member of a static PortChannel only if the following configurations are the same in the port and the PortChannel:

- Speed
- Mode
- Rate mode
- Port VSAN
- Trunking mode
- Allowed VSAN list or VF-ID list

After the members are added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Generation 1 PortChannel Limitations" section on page 6-14 and "Graceful Shutdown" section on page 2-9).

Compatibility Check

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a PortChannel. The compatibility check is performed before a port is added to the PortChannel.

The check ensures that the following parameters and settings match at both ends of a PortChannel:

- Capability parameters (type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends).
- Administrative compatibility parameters (speed, mode, rate mode, port VSAN, allowed VSAN list, and port security).

  Note Ports in shared rate mode cannot form a PortChannel or a trunking PortChannel.

- Operational parameters (remote switch WWN and trunking mode).

A port addition procedure fails if the capability and administrative parameters in the remote switch are incompatible with the capability and administrative parameters in the local switch. If the compatibility check is successful, the interfaces are operational and the corresponding compatibility parameter settings apply to these interfaces.

Suspended and Isolated States

If the operational parameters are incompatible, the compatibility check fails and the interface is placed in a suspended or isolated state based on the configured mode:

- An interface enters the suspended state if the interface is configured in the ON mode.
- An interface enters the isolated state if the interface is configured in the ACTIVE mode.

Forcing an Interface Addition

You can force the port configuration to be overwritten by the PortChannel. In this case, the interface is added to a PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the addition.

Note When PortChannels are created from within an interface, the **force** option cannot be used.

After the members are forcefully added, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Graceful Shutdown" section on page 2-9).

Interface Deletion from a PortChannel

When a physical interface is deleted from the PortChannel, the channel membership is automatically updated. If the deleted interface is the last operational interface, then the PortChannel status is changed to a down state. Deleting an interface from a PortChannel decreases the channel size and bandwidth of the PortChannel.

- If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down. You must explicitly enable those ports again.
- If you use the ACTIVE mode, then the PortChannel ports automatically recover from the deletion.

After the members are deleted, regardless of the mode (ACTIVE and ON) used, the ports at either end are gracefully brought down, indicating that no frames are lost when the interface is going down (see the "Generation 1 PortChannel Limitations" section on page 6-14 and "Graceful Shutdown" section on page 2-9).

PortChannel Protocols

In earlier Cisco SAN-OS releases, PortChannels required additional administrative tasks to support synchronization. The Cisco NX-OS software provides robust error detection and synchronization capabilities. You can manually configure channel groups or they can be automatically created. In both cases, the channel groups have the same capability and configurational parameters. Any change in configuration applied to the associated PortChannel interface is propagated to all members of the channel group.

A protocol to exchange PortChannel configurations is available in all Cisco MDS switches. This addition simplifies PortChannel management with incompatible ISLs. An additional autocreation mode enables ISLs with compatible parameters to automatically form channel groups without manual intervention.

The PortChannel protocol is enabled by default.

The PortChannel protocol expands the PortChannel functional model in Cisco MDS switches. It uses the exchange peer parameters (EPP) services to communicate across peer ports in an ISL. Each switch uses the information received from the peer ports along with its local configuration and operational values to decide if it should be part of a PortChannel. The protocol ensures that a set of ports are eligible to be part of the same PortChannel. They are only eligible to be part of the same PortChannel if all the ports have a compatible partner.

The PortChannel protocol uses two subprotocols:

- Bringup protocol—Automatically detects misconfigurations so you can correct them. This protocol synchronizes the PortChannel at both ends so that all frames for a given flow (as identified by the source FC ID, destination FC ID and OX_ID) are carried over the same physical link in both directions. This helps make applications such as write acceleration, work for PortChannels over FCIP links.
- Autocreation protocol—Automatically aggregates compatible ports into a PortChannel.

This section describes how to configure the PortChannel protocol and includes the following sections:

- Channel Group Creation, page 6-10
- Autocreation, page 6-11
- Manually Configured Channel Groups, page 6-12

Channel Group Creation

Note Channel groups are not supported on internal ports in the Cisco Fabric Switch for HP c-Class BladeSystem and the Cisco Fabric Switch for IBM BladeSystem.

Assuming link A1-B1 comes up first (see Figure 6-6), that link is operational as an individual link.When the next link comes up, for example, A2-B2, the PortChannel protocol identifies if this link is compatible with link A1-B1 and automatically creates channel groups 10 and 20 in the respective switches. If link A3-B3 can join the channel groups (the PortChannels), the respective ports have compatible configurations. If link A4-B4 operates as an individual link, it is because of the incompatible configuration of the two end ports with the other member ports in this channel group.

Figure 6-6 Autocreating Channel Groups

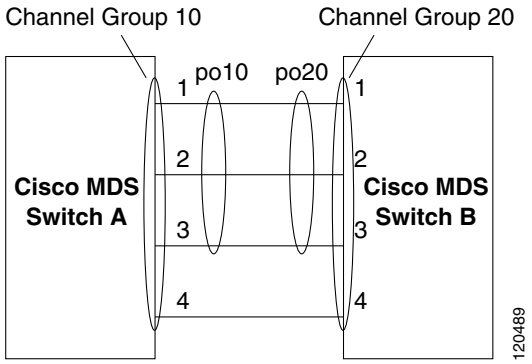

The channel group numbers are selected dynamically, and as such, the administrative configuration of the ports forming the channel group at either end are applicable to the newly created channel group. The channel group number being chosen dynamically may be different across reboots for the same set of PortChannels based on the order of ports that are initialized in the switch.

Table 6-2 identifies the differences between user-configured and auto-configured channel groups.

Table 6-2 Channel Group Configuration Differences

User-Configured Channel Group	Autocreated Channel Group
Manually configured by the user.	Created automatically when compatible links come up between two compatible switches, if channel group autocreation is enabled in all ports at both ends.
Member ports cannot participate in autocreation of channel groups. The autocreation feature cannot be configured.	None of these ports are members of a user-configured channel group.
You can form the PortChannel with a subset of the ports in the channel group. Incompatible ports remain in a suspended or isolated state depending on the ON or ACTIVE mode configuration.	All ports included in the channel group participate in the PortChannel—no member port becomes isolated or suspended; instead, the member port is removed from the channel group when the link is found to be incompatible.
Any administrative configuration made to the PortChannel is applied to all ports in the channel group, and you can save the configuration for the PortChannel interface.	Any administrative configuration made to the PortChannel is applied to all ports in the channel group, but the configurations are saved for the member ports; no configuration is saved for the PortChannel interface. You can explicitly convert this channel group, if required.
You can remove any channel group and add members to a channel group.	You cannot remove a channel group, or add/remove any of its members. The channel group is removed when no member ports exist.

Note Autocreation is not supported as of MDS NX-OS Release 4.1(1b) and later.

Autocreation

The autocreation protocol has the following functionality:

- A port is not allowed to be configured as part of a PortChannel when the autocreation feature is enabled. These two configurations are mutually exclusive.
- Autocreation must be enabled in both the local and peer ports to negotiate a PortChannel.
- Aggregation occurs in one of two ways:
  - A port is aggregated into a compatible autocreated PortChannel.
  - A port is aggregated with another compatible port to form a new PortChannel.

- Newly created PortChannels are allocated from the maximum possible PortChannel (128 for Generation 1 or a combination of Generation 1 and Generation 2 switches, or 256 for Generation 2 switches) in a decreasing order based on availability. If all 128 (or 256) numbers are used up, aggregation is not allowed.
- You cannot change the membership or delete an autocreated PortChannel.
- When you disable autocreation, all member ports are removed from the autocreated PortChannel.
- Once the last member is removed from an autocreated PortChannel, the channel is automatically deleted and the number is released for reuse.
- An autocreated PortChannel is not persistent through a reboot. An autocreated PortChannel can be manually configured to appear the same as a persistent PortChannel. Once the PortChannel is made persistent, the autocreation feature is disabled in all member ports.
- You can enable or disable the autocreation feature on a per-port basis or for all ports in the switch. When this configuration is enabled, the channel group mode is assumed to be active. The default for this task is disabled.
- If autocreation of channel groups is enabled for an interface, you must first disable autocreation before downgrading to earlier software versions or before configuring the interface in a manually configured channel group.

Tip When enabling autocreation in any switch in the Cisco MDS 9000 Family, we recommend that you retain at least one interconnected port between the switches without any autocreation configuration. If all ports between two switches are configured with the autocreation feature at the same time, you may face a possible traffic disruption between these two switches as the ports are automatically disabled and reenabled when ports are added to an autocreated PortChannel.

Manually Configured Channel Groups

A user-configured channel group cannot be converted to an autocreated channel group. However, you can convert an autocreated channel group to a manual channel group. Once performed, this task is irreversible. The channel group number does not change, but the member ports operate according to the properties of the manually configured channel group, and the autocreation of channel group is implicitly disabled for all member ports.

Tip If you enable persistence, be sure to enable it at both ends of the PortChannel.

Prerequisites for PortChannels

Before configuring a PortChannel, consider the following guidelines:

- Configure the PortChannel across switching modules to implement redundancy on switching module reboots or upgrades.
- Ensure that one PortChannel is not connected to different sets of switches. PortChannels require point-to-point connections between the same set of switches.

Note On switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, you can configure a maximum of 128 PortChannels. On switches with only Generation 2 switching modules, or Generation 2 and Generation 3 switching modules, you can configure a maximum of 256 PortChannels.

If you misconfigure PortChannels, you may receive a misconfiguration message. If you receive this message, the PortChannel's physical links are disabled because an error has been detected.

A PortChannel error is detected if the following requirements are not met:

- Each switch on either side of a PortChannel must be connected to the same number of interfaces.
- Each interface must be connected to a corresponding interface on the other side (see Figure 6-8 for an example of an invalid configuration).
- Links in a PortChannel cannot be changed after the PortChannel is configured. If you change the links after the PortChannel is configured, be sure to reconnect the links to interfaces within the PortChannel and reenable the links.

If all three conditions are not met, the faulty link is disabled.

Enter the **show interface** command for that interface to verify that the PortChannel is functioning as required.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

- General Guidelines and Limitations, page 6-13
- Generation 1 PortChannel Limitations, page 6-14
- F and TF PortChannel Limitations, page 6-14
- Valid and Invalid PortChannel Examples, page 6-15

General Guidelines and Limitations

Cisco MDS 9000 Family switches support the following number of PortChannels per switch:

- Switches with only Generation 1 switching modules do not support F and TF PortChannels.
- Switches with Generation 1 switching modules, or a combination of Generation 1 and Generation 2 switching modules, support a maximum of 128 PortChannels. Only Generation 2 ports can be included in the PortChannels.
- Switches with only Generation 2 switching modules or Generation 2 and Generation 3 modules support a maximum of 256 PortChannels with 16 interfaces per PortChannel.
- A PortChannel number refers to the unique identifier for each channel group. This number ranges from of 1 to 256.

Generation 1 PortChannel Limitations

This section includes the restrictions on creation and addition of PortChannel members to a PortChannel on Generation 1 hardware:

- The 32-port 2-Gbps or 1-Gbps switching module.
- The MDS 9140 and 9120 switches.

When configuring the host-optimized ports on Generation 1 hardware, the following PortChannel guidelines apply:

- If you execute the **write erase** command on a 32-port switching module, and then copy a saved configuration to the switch from a text file that contains the **no system default switchport shutdown** command, you need to copy the text file to the switch again for the E ports to come up without manual configuration.
- Any (or all) full line rate port(s) in the Cisco MDS 9100 Series can be included in a PortChannel.
- The host-optimized ports in the Cisco MDS 9100 Series are subject to the same PortChannel rules as 32-port switching modules; only the first port of each group of 4 ports is included in a PortChannel.
  - You can configure only the first port in each 4-port group as an E port (for example, the first port in ports 1–4, the fifth port in ports 5–8, and so on). If the first port in the group is configured as a PortChannel, the other three ports in each group (ports 2–4, 6–8, and so on) are not usable and remain in the shutdown state.
  - If any of the other three ports are configured in a no shutdown state, you cannot configure the first port to be a PortChannel. The other three ports continue to remain in a no shutdown state.

F and TF PortChannel Limitations

The following guidelines and restrictions are applicable for F and TF PortChannels:

- The ports must be in F mode.
- Automatic creation is not supported.
- The PortChannel interface must be in ACTIVE mode when multiple FCIP interfaces are grouped with WA.
- ON mode is not supported. Only ACTIVE-ACTIVE mode is supported. By default, the mode is ACTIVE on the NPV switches.
- Devices logged in through F PortChannel on an MDS switch are not supported in IVR non-NAT configuration. The devices are supported only in IVR NAT configuration.
- Port security rules are enforced only on physical pWWNs at the single link level.
- FC-SP authenticates only the first physical FLOGI of every PortChannel member.
- Since the FLOGI payload carries only the VF bits to trigger the use of a protocol after the FLOGI exchange, those bits will be overridden. In the case of the NPV switches, the core has a Cisco WWN and will try to initiate the PCP protocol.
- The name server registration of the N ports logging in through an F PortChannel will use the fWWN of the PortChannel interface.
- DPVM configuration is not supported.
- The PortChannel port VSAN cannot be configured using DPVM.

- The Dynamic Port VSAN Management (DPVM) database will be queried only for the first physical FLOGI of each member, so that the port VSAN can be configured automatically.
- DPVM does not bind FC_IDs to VSANs, but pWWNs to VSANs. It will be queried only for the physical FLOGI.

Valid and Invalid PortChannel Examples

PortChannels are created with default values. You can change the default configuration just like any other physical interface.

Figure 6-7 provides examples of valid PortChannel configurations.

Figure 6-7 Valid PortChannel Configurations

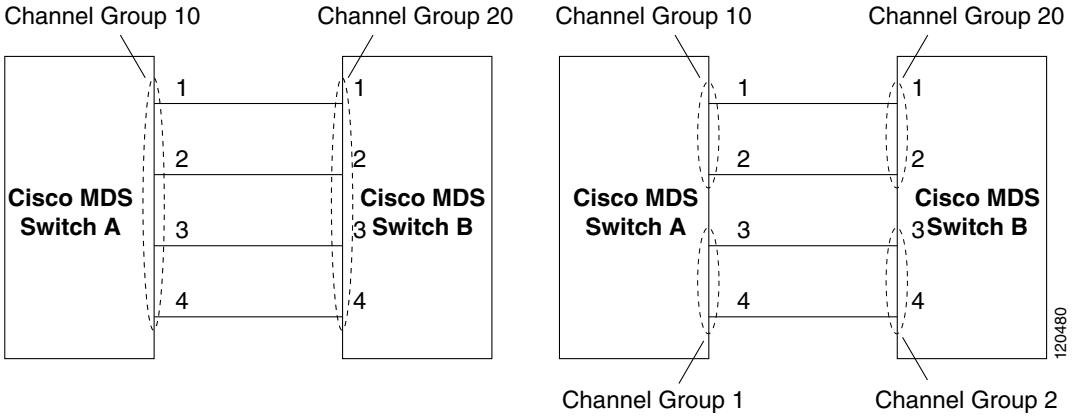

Figure 6-8 provides examples of invalid configurations. Assuming that the links are brought up in the 1, 2, 3, 4 sequence, links 3 and 4 will be operationally down as the fabric is misconfigured.

Figure 6-8 Misconfigured Configurations

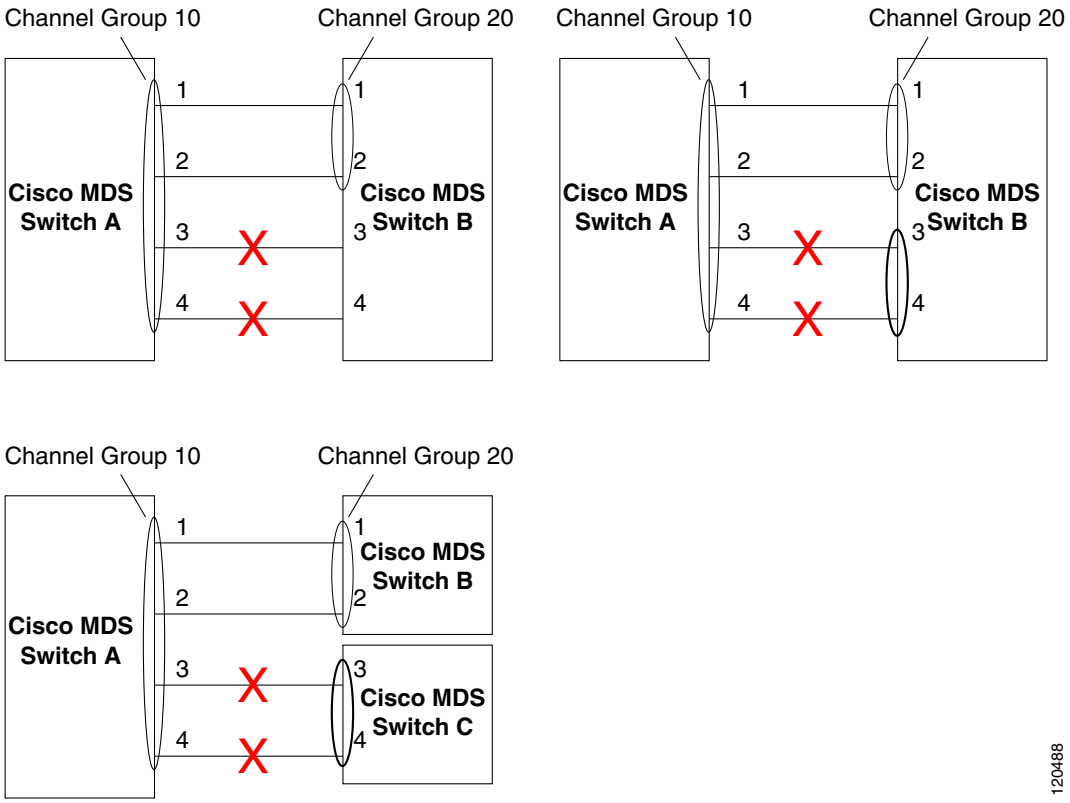

Default Settings

Table 6-3 lists the default settings for PortChannels.

Table 6-3 Default PortChannel Parameters

Parameters	Default
PortChannels	FSPF is enabled by default.
Create PortChannel	Administratively up.
Default PortChannel mode	ON mode on non-NPV and NPIV core switches. ACTIVE mode on NPV switches.
Autocreation	Disabled.

Configuring PortChannels

This section includes the following topics:

- Configuring PortChannels Using the WizardCreating a PortChannel, page 6-17Configuring PortChannels Using the WizardCreating a PortChannel, page 6-17
- Configuring the PortChannel Mode, page 6-17
- Deleting PortChannels, page 6-18
- Adding an Interface to a PortChannel, page 6-18
- Forcing an Interface Addition, page 6-19
- Deleting an Interface from a PortChannel, page 6-19
- Enabling and Configuring Autocreation, page 6-20
- Converting to Manually Configured Channel Groups, page 6-20

Configuring PortChannels Using the WizardCreating a PortChannel

Detailed Steps

To create a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface port-channel 1` `switch(config-if)#`	Configures the specified PortChannel (1) using the default ON mode.

Configuring the PortChannel Mode

By default, the CLI and the Device Manager create the PortChannel in ON mode in the NPIV core switches and ACTIVE mode on the NPV switches. DCNM-SAN creates all PortChannels in ACTIVE mode. We recommend that you create PortChannels in ACTIVE mode.

Restrictions

- An F PortChannel is supported only on ACTIVE mode.

Detailed Steps

To configure ACTIVE mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface port-channel 1` `switch(config-if)#`	Configures the specified PortChannel (1) using the default ON mode.

	Command	Purpose
Step 3	`switch(config-if)# channel mode active`	Configures the ACTIVE mode.
	`switch(config-if)# no channel mode active`	Reverts to the default ON mode.

Deleting PortChannels

Detailed Steps

To delete a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

Step 2	`switch(config)# no interface port-channel 1`
	`port-channel 1 deleted and all its members disabled`
	`please do the same operation on the switch at the other end of the port-channel`
	`switch(config)#`
		Deletes the specified PortChannel (1), its associated interface mappings, and the hardware associations for this PortChannel.

Adding an Interface to a PortChannel

Detailed Steps

To add an interface to a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/15` `switch(config-if)#`	Configures the specified port interface (fc1/15).
Step 3	`switch(config-if)# channel-group 15`	Adds physical Fibre Channel port 1/15 to channel group 15. If channel group 15 does not exist, it is created. The port is shut down.

To add a range of ports to a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.

	Command	Purpose
Step 2	`switch(config)# interface fc1/1 - 5` `switch(config-if)#`	Configures the specified range of interfaces. In this example, interfaces from 1/1 to 1/5 are configured.
Step 3	`switch(config-if)# channel-group 2`	Adds physical interfaces 1/1, 1/2, 1/3, 1/4, and 1/5 to channel group 2. If channel group 2 does not exist, it is created. If the compatibility check is successful, the interfaces are operational and the corresponding states apply to these interfaces.

Note By default, the CLI adds a interface normally to a PortChannel, while DCNM-SAN adds the interface by force, unless specified explicitly.

Forcing an Interface Addition

Detailed Steps

To force the addition of a port to a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Specifies the interface fc1/1.
Step 3	`switch(config-if)# channel-group 1 force`	Forces the addition of the physical port for interface fc1/1 to channel group 1. The port is shut down.

Deleting an Interface from a PortChannel

Detailed Steps

To delete a physical interface (or a range of physical interfaces) from a PortChannel, follow these steps:

	Command	Purpose
Step 1	`switch(config)# interface fc1/1` `switch(config-if)#`	Enters the selected physical interface level.
	`switch(config)# interface fc1/1 - 5` `switch(config-if)#`	Enters the selected range of physical interfaces.
Step 2	`switch(config-if)# no channel-group 2` `switch(config-if)#`	Deletes the physical Fibre Channel interfaces in channel group 2.

Enabling and Configuring Autocreation

Detailed Steps

To configure automatic channel groups, follow these steps:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/13` `switch(config- if)#`	Enters the configuration mode for the selected interface(s).
Step 3	`switch(config- if)# channel-group auto`	Automatically creates the channel group for the selected interface(s).
	`switch(config- if)# no channel-group auto`	Disables the autocreation of channel groups for this interface, even if the system default configuration may have autocreation enabled.

Converting to Manually Configured Channel Groups

You can convert autocreated channel group to a user-configured channel group using the **port-channel** *channel-group-number* **persistent** EXEC command. If the PortChannel does not exist, this command is not executed.

Verifying PortChannel Configuration

To display PortChannel configuration information, perform one of the following tasks:

Command	Purpose
show port-channel summary	Displays a summary of PortChannels within the switch. A one-line summary of each PortChannel provides the administrative state, the operational state, the number of attached and active interfaces (up), and the first operational port (FOP), which is the primary operational interface selected in the PortChannel to carry control-plane traffic (no load-balancing). The FOP is the first port that comes up in a PortChannel and can change if the port goes down. The FOP is also identified by an asterisk (*).
show port-channel database	Displays the PortChannel configured in the default ON mode and ACTIVE mode.
show port-channel consistency	Displays the consistency status without details.
show port-channel consistency detail	Displays the consistency status with details.
show port-channel usage	Displays the PortChannel usage.
show port-channel compatibility-parameters	Displays the PortChannel compatibility.

Command	Purpose
show interface fc *slot/port*	Displays autocreated PortChannels.
show port-channel database interface port-channel *number*	Displays the specified PortChannel interface.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

You can view specific information about existing PortChannels at any time from EXEC mode. The following **show** commands provide further details on existing PortChannels. You can force all screen output to go to a printer or save it to a file. See Examples 6-1 to 6-6.

Example 6-1 Displays the PortChannel Summary

```
switch# show port-channel summary
--------------------------------------------------------------------------------
Interface                 Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
port-channel 77                2                 0                 --
port-channel 78                2                 0                 --
port-channel 79                2                 2                 fcip200
```

Example 6-2 Displays the PortChannel Configured in the Default ON Mode

```
switch# show port-channel database
port-channel 77
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
             fcip2    [down]
port-channel 78
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is on
    Operational channel mode is on
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

Example 6-3 Displays the PortChannel Configured in the ACTIVE Mode

```
switch# show port-channel database
port-channel 77
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fcip1    [down]
```

```
            fcip2    [down]
port-channel 78
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    2 ports in total, 0 ports up
    Ports:   fc2/1    [down]
             fc2/5    [down]
port-channel 79
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    First operational port is fcip200
    2 ports in total, 2 ports up
    Ports:   fcip101  [up]
             fcip200  [up] *
```

The **show port-channel consistency** command has two options: without details and with details.

Example 6-4 Displays the Consistency Status without Details

```
switch# show port-channel consistency
Database is consistent
```

Example 6-5 Displays the Consistency Status with Details

```
switch# show port-channel consistency detail
Authoritative port-channel database:
================================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
================================================
database 1: from module 5
================================================
totally 3 port-channels
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
================================================
database 2: from module 4
================================================
totally 3 port-channels
```

```
port-channel 77:
    2 ports, first operational port is none
    fcip1    [down]
    fcip2    [down]
port-channel 78:
    2 ports, first operational port is none
    fc2/1    [down]
    fc2/5    [down]
port-channel 79:
    2 ports, first operational port is fcip200
    fcip101  [up]
    fcip200  [up]
...
```

The **show port-channel usage** command displays details of the used and unused PortChannel numbers.

Example 6-6 Displays the PortChannel Usage

```
switch# show port-channel usage
Totally 3 port-channel numbers used
===================================
Used  :   77 - 79
Unused:   1 - 76 , 80 - 256
```

Example 6-7 Displays the PortChannel Compatibility

```
switch# show port-channel compatibility-parameters
physical port layer                 fibre channel or ethernet
    port mode                           E/AUTO only
    trunk mode
    speed
    port VSAN
    port allowed VSAN list
```

Use the existing **show** commands to obtain further details on autocreated channel group attributes. Autocreated PortChannels are indicated explicitly to help differentiate them from the manually created PortChannels. See Examples 6-8 to 6-10.

Example 6-8 Displays Autocreated PortChannels

```
switch# show interface fc1/1
fc1/1 is trunking
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:0a:00:0b:5f:3b:fe:80
    ...
    Receive data field Size is 2112
    Beacon is turned off
    Port-channel auto creation is enabled
    Belongs to port-channel 123
...
```

Example 6-9 Displays the Specified PortChannel Interface

```
switch# show port-channel database interface port-channel 128
port-channel 128
    Administrative channel mode is active
    Operational channel mode is active
    Last membership update succeeded
    Channel is auto created
    First operational port is fc1/1
```

```
    1 ports in total, 1 ports up
    Ports:   fc1/1   [up] *
```

Example 6-10 Displays the PortChannel Summary

```
switch# show port-channel summary
--------------------------------------------------------------------------------
Interface                Total Ports        Oper Ports        First Oper Port
--------------------------------------------------------------------------------
port-channel 1                1                0                 --
port-channel 2                1                1                 fc8/13
port-channel 3                0                0                 --
port-channel 4                0                0                 --
port-channel 5                1                1                 fc8/3
port-channel 6                0                0                 --
```

Configuration Examples for F and TF PortChannels

This example shows how to configure F PortChannel in shared mode and bring up the link (not supported on the MDS 91x4 switches) between F ports on the NPIV core switches and NP ports on the NPV switches:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch.

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

```
switch(config)# feature npiv
```

Step 3 Create the PortChannel on the MDS core switch:

```
switch(config)# interface port-channel 1
switch(config-if)# switchport mode F
switch(config-if)# channel mode active
switch(config-if)# switchport trunk mode off
switch(config-if)# switchport rate-mode shared
switch(config-if)# exit
```

Step 4 Configure the PortChannel member interfaces on the core switch:

```
switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport trunk mode off
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# channel-group 1
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the PortChannel on the NPV switch:

```
switch(config)# interface port-channel 1
switch(config-if)# switchport mode NP
switch(config-if)# switchport rate-mode shared
switch(config-if)# exit
```

Step 6 Configure the PortChannel member interfaces on the NPV switch:

```
switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# switchport mode NP
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode shared
switch(config-if)# switchport trunk mode off
switch(config-if)# channel-group 1
switch(config-if)# no shut
switch(config-if)# exit
```

Step 7 Set the administrative state of all the PortChannel member interfaces in both NPIV core switch and the NPV switch to ON:

```
switch(config)# interface fc1/1-3
switch(config-if)# shut
switch(config-if)# no shut

switch(config)# interface fc2/1-3
switch(config-if)# shut
switch(config-if)# no shut
```

Note The speed configuration must be the same for all member interfaces in a PortChannel. While configuring the channel in dedicated mode, ensure that required bandwidth is available to the ports.

This example shows how to configure channeling in dedicated mode and bring up the TF-TNP PortChannel link between TF ports in the NPIV core switch, and TNP ports in the NPV switch:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

```
switch(config)# feature npiv
```

Step 3 Create the PortChannel on the MDS core switch:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport mode F
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# channel mode active
switch(config-if)# exit
```

Step 4 Configure the PortChannel member interfaces on the MDS core switch in dedicated mode:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the PortChannel in dedicated mode on the NPV switch:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport rate-mode dedicated
```

```
switch(config-if)# switchport mode NP
switch(config-if)# no shut
switch(config-if)# exit
```

Step 6 Configure the PortChannel member interfaces on the NPV switch in dedicated mode:

```
switch(config)# interface fc3/1-3
switch(config-if)# shut
switch(config-if)# switchport mode NP
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Step 7 Set the administrative state of all the PortChannel member interfaces in both NPIV core switch and the NPV switch to ON:

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# no shut

switch(config)# interface fc3/1-3
switch(config-if)# shut
switch(config-if)# no shut
```

Configuring N Port Virtualization

- Information About N Port Virtualization, page 7-1
- Guidelines and Limitations, page 7-8
- Configuring N Port Virtualization, page 7-10
- Verifying NPV Configuration, page 7-13

Information About N Port Virtualization

This section includes the following topics:

- NPV Overview, page 7-1
- N Port Identifier Virtualization, page 7-2
- N Port Virtualization, page 7-2
- NPV Mode, page 7-4
- NP Ports, page 7-5
- NP Links, page 7-5
- Default Port Numbers, page 7-6
- NPV CFS Distribution over IP, page 7-7
- NPV Traffic Management, page 7-7
- Multiple VSAN Support, page 7-8

NPV Overview

N port virtualization (NPV) reduces the number of Fibre Channel domain IDs in SANs. Switches operating in the NPV mode do not join a fabric. They pass traffic between NPV core switch links and end devices, which eliminates the domain IDs for these edge switches.

NPV is supported by the following Cisco MDS 9000 switches and Cisco Nexus 5000 Series switches only:

- Cisco MDS 9124 Multilayer Fabric Switch
- Cisco MDS 9134 Fabric Switch
- Cisco MDS 9148 Multilayer Fabric Switch

- Cisco Fabric Switch for HP c-Class BladeSystem
- Cisco Fabric Switch for IBM BladeCenter
- Cisco Nexus 5000 Series switches

Note NPV is available on these switches only while in NPV mode; if in switch mode, NPV is not available.

N Port Identifier Virtualization

N port identifier virtualization (NPIV) provides a means to assign multiple FC IDs to a single N port. This feature allows multiple applications on the N port to use different identifiers and allows access control, zoning, and port security to be implemented at the application level. Figure 7-1 shows an example application using NPIV.

Figure 7-1 NPIV Example

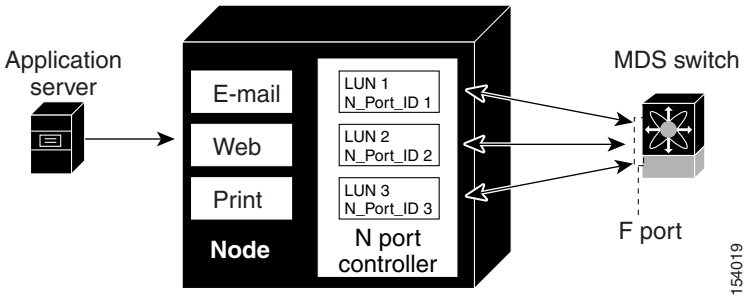

You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

Note All of the N port identifiers are allocated in the same VSAN.

N Port Virtualization

Typically, Fibre Channel networks are deployed using a core-edge model with a large number of fabric switches connected to edge devices. Such a model is cost-effective because the per port cost for director class switches is much higher than that of fabric switches. However, as the number of ports in the fabric increases, the number of switches deployed also increases, and you can end up with a significant increase in the number of domain IDs (the maximum number supported is 239). This challenge becomes even more difficult when additional blade chassis are deployed in Fibre Channel networks.

NPV addresses the increase in the number of domain IDs needed to deploy a large number of the ports by making a fabric or blade switch appear as a host to the core Fibre Channel switch, and as a Fibre Channel switch to the servers in the fabric or blade switch. NPV aggregates multiple locally connected N ports into one or more external NP links, which shares the domain ID of the NPV core switch among multiple NPV switches. NPV also allows multiple devices to attach to same port on the NPV core switch, which reduces the need for more ports on the core.

Figure 7-2 Cisco NPV Fabric Configuration

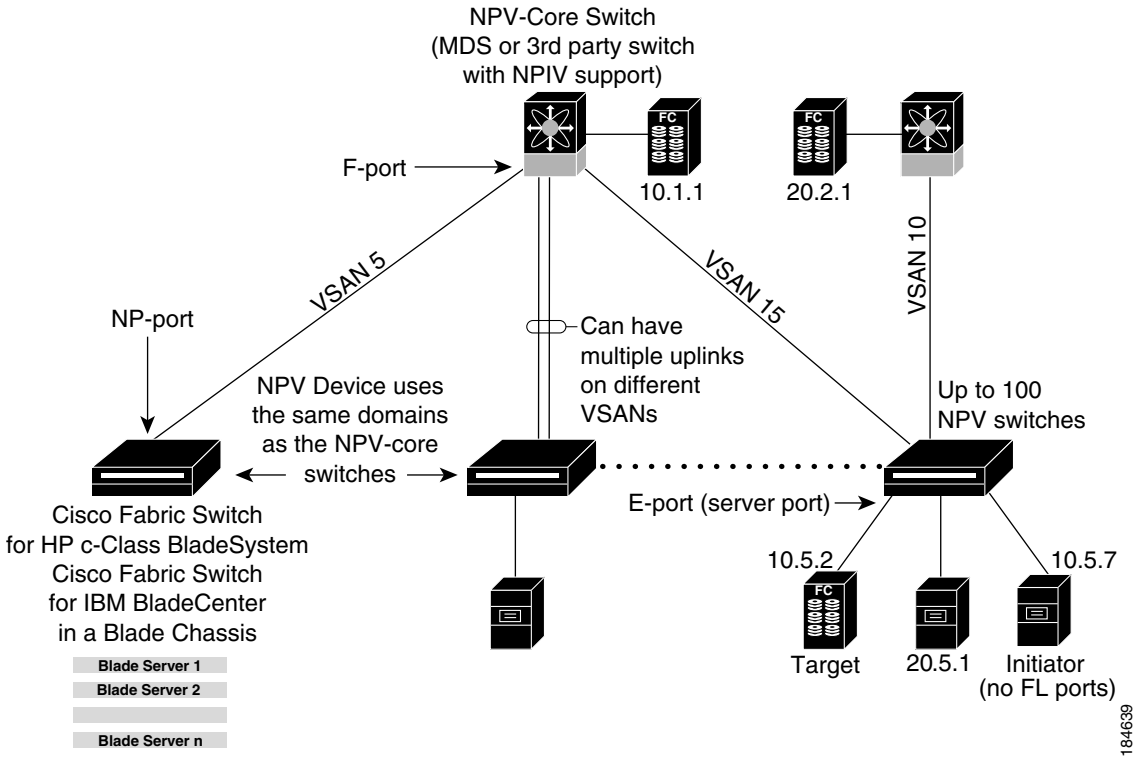

While NPV is similar to N port identifier virtualization (NPIV), it does not offer exactly the same functionality. NPIV provides a means to assign multiple FC IDs to a single N port, and allows multiple applications on the N port to use different identifiers. NPIV also allows access control, zoning, and port security to be implemented at the application level. NPV makes use of NPIV to get multiple FCIDs allocated from the core switch on the NP port.

Figure 7-3 shows a more granular view of an NPV configuration at the interface level.

Figure 7-3 Cisco NPV Configuration–Interface View

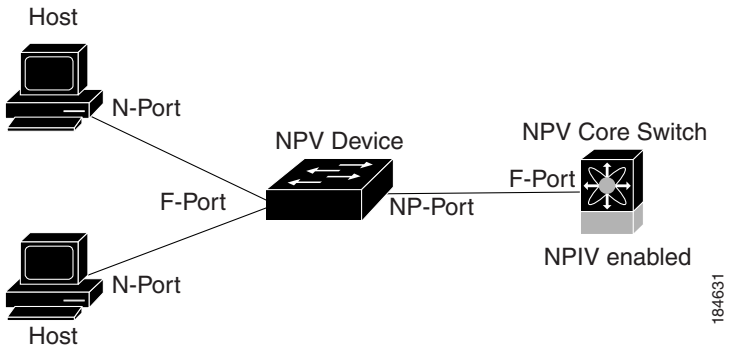

NPV Mode

A switch is in NPV mode after a user has enabled NPV and the switch has successfully rebooted. NPV mode applies to an entire switch. All end devices connected to a switch that is in NPV mode must log in as an N port to use this feature (loop-attached devices are not supported). All links from the edge switches (in NPV mode) to the NPV core switches are established as NP ports (not E ports), which are used for typical interswitch links. NPIV is used by the switches in NPV mode to log in to multiple end devices that share a link to the NPV core switch.

Note In-order data delivery is not required in NPV mode because the exchange between two end devices always takes the same uplink to the core from the NPV device. For traffic beyond the NPV device, core switches will enforce in-order delivery if needed and/or configured.

After entering NPV mode, only the following commands are available:

```
aaa             Configure aaa functions
arp             [no] remove an entry from the ARP cache
banner          Configure banner message
boot            Configure boot variables
callhome        Enter the callhome configuration mode
cli             CLI configuration commands
clock           Configure time-of-day clock
do              EXEC command
end             Exit from configure mode
exit            Exit from configure mode
fcanalyzer      Configure cisco fabric analyzer
fcrxbbcredit    Enable extended rx b2b credit configuration
fips            Enable/Disable FIPS mode
hw-module       Enable/Disable OBFL information
interface       Select an interface to configure
ip              Configure IP features
ipv6            Configure IPv6 features
line            Configure a terminal line
logging         Modify message logging facilities
no              Negate a command or set its defaults
npv             Config commands for FC N_port Virtualizer
ntp             NTP Configuration
port-track      Configure Switch port track config
power           Configure power supply
poweroff        Poweroff a module in the switch
radius          Configure RADIUS configuration
radius-server   Configure RADIUS related parameters
rate-mode       Configure rate mode oversubscription limit
rmon            Remote Monitoring
role            Configure roles
snmp-server     Configure snmp server
ssh             Configure SSH parameters
switchname      Configure system's network name
system          System config command
tacacs+         Enable tacacs+
telnet          Enable telnet
username        Configure user information.
wwn             Set secondary base MAC addr and range for additional WWNs
```

NP Ports

An NP port (proxy N port) is a port on a device that is in NPV mode and connected to the NPV core switch using an F port. NP ports behave like N ports except that in addition to providing N port behavior, they also function as proxies for multiple, physical N ports.

NP Links

An NP link is basically an NPIV uplink to a specific end device. NP links are established when the uplink to the NPV core switch comes up; the links are terminated when the uplink goes down. Once the uplink is established, the NPV switch performs an internal FLOGI to the NPV core switch, and then (if the FLOGI is successful) registers itself with the NPV core switch's name server. Subsequent FLOGIs from end devices in this NP link are converted to FDISCs. For more details refer to the "Internal FLOGI Parameters" section on page 7-5.

Server links are uniformly distributed across the NP links. All the end devices behind a server link will be mapped to only one NP link.

Internal FLOGI Parameters

When an NP port comes up, the NPV device first logs itself in to the NPV core switch and sends a FLOGI request that includes the following parameters:

- The fWWN (fabric port WWN) of the NP port used as the pWWN in the internal login.
- The VSAN-based sWWN (switch WWN) of the NPV device used as nWWN (node WWN) in the internal FLOGI.

After completing its FLOGI request, the NPV device registers itself with the fabric name server using the following additional parameters:

- Switch name and interface name (for example, fc1/4) of the NP port is embedded in the symbolic port name in the name server registration of the NPV device itself.
- The IP address of the NPV device is registered as the IP address in the name server registration of the NPV device.

Note The BB_SCN of internal FLOGIs on NP ports is always set to zero. The BB_SCN is supported at the F-port of the NPV device.

Figure 7-4 shows the internal FLOGI flows between an NPV core switch and an NPV device.

Figure 7-4 Internal FLOGI Flows

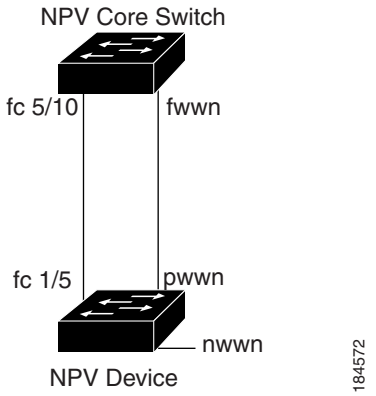

Table 7-1 identifies the internal FLOGI parameters that appear in Figure 7-4.

Table 7-1 Internal FLOGI Parameters

Parameter	Derived From
pWWN	The fWWN of the NP port.
nWWN	The VSAN-based sWWN of the NPV device.
fWWN	The fWWN of the F port on the NPV core switch.
symbolic port name	The switch name and NP port interface string. **Note** If there is no switch name available, then the output will display "switch." For example, switch: fc1/5.
IP address	The IP address of the NPV device.
symbolic node name	The NPV switch name.

Although fWWN-based zoning is supported for NPV devices, it is not recommended because:

- Zoning is not enforced at the NPV device (rather, it is enforced on the NPV core switch).
- Multiple devices behind an NPV device log in via the same F port on the core (they use same fWWN and cannot be separated into different zones).
- The same device might log in using different fWWNs on the core switch (depending on the NPV link it uses) and may need to be zoned using different fWWNs.

Default Port Numbers

Port numbers on NPV-enabled switches will vary depending on the switch model. For details about port numbers for NPV-eligible switches, see the *Cisco NX-OS Family Licensing Guide*.

NPV CFS Distribution over IP

NPV devices use only IP as the transport medium. CFS uses multicast forwarding for CFS distribution. NPV devices do not have ISL connectivity and FC domain. To use CFS over IP, multicast forwarding has to be enabled on the Ethernet IP switches all along the network that physically connects the NPV switch. You can also manually configure the static IP peers for CFS distribution over IP on NPV-enabled switches. For more information, see the *Cisco MDS 9000 Family NX-OS System Management Configuration Guide*.

NPV Traffic Management

This sections discusses the following aspects of load balancing:

- Auto, page 7-7
- Traffic Map, page 7-7
- Disruptive, page 7-8

Auto

Before Cisco MDS SAN-OS Release 3.3(1a), NPV supported automatic selection of external links. When a server interface is brought up, an external interface with the minimum load is selected from the available links. There is no manual selection on the server interfaces using the external links. Also, when a new external interface was brought up, the existing load was not distributed automatically to the newly available external interface. This newly brought up interface is used only by the server interfaces that come up after this interface.

Traffic Map

As in Cisco MDS SAN-OS Release 3.3(1a) and NX-OS Release 4.1(1a), NPV supports traffic management by allowing you to select and configure the external interfaces that the server uses to connect to the core switches.

Note When the NPV traffic management is configured, the server uses only the configured external interfaces. Any other available external interface will not be used.

The NPV traffic management feature provides the following benefits:

- Facilitates traffic engineering by providing dedicated external interfaces for the servers connected to NPV.
- Uses the shortest path by selecting external interfaces per server interface.
- Uses the persistent FC ID feature by providing the same traffic path after a link break, or reboot of the NPV or core switch.
- Balances the load by allowing the user to evenly distribute the load across external interfaces.

Disruptive

Disruptive load balance works independent of automatic selection of interfaces and a configured traffic map of external interfaces. This feature forces reinitialization of the server interfaces to achieve load balance when this feature is enabled and whenever a new external interface comes up. To avoid flapping the server interfaces too often, enable this feature once and then disable it whenever the needed load balance is achieved.

If disruptive load balance is not enabled, you need to manually flap the server interface to move some of the load to a new external interface.

Multiple VSAN Support

By grouping devices into different NPV sessions based on VSANs, it is possible to support multiple VSANs on the NPV-enabled switch. The correct uplink must be selected based on the VSAN that the uplink is carrying.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature:

- NPV Guidelines and Requirements, page 7-8
- NPV Traffic Management Guidelines, page 7-9
- DPVM Configuration Guidelines, page 7-9
- NPV and Port Security Configuration Guidelines, page 7-10

NPV Guidelines and Requirements

Following are recommended guidelines and requirements when deploying NPV:

- NPV core switches must support NPIV.
- You can have up to 100 NPV devices.
- Nondisruptive upgrades are supported. See the *Cisco MDS 9000 Family NX-OS Fundamentals Configuration Guide*.
- Port tracking is supported. See the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.
- You can configure zoning for end devices that are connected to NPV devices using all available member types on the NPV core switch. If fWWN, sWWN, domain, or port-based zoning is used, then fWWN, sWWN or the domain/port of the NPV core switch should be used.
- Port security is supported on the NPV core switch for devices logged in via NPV.
- NPV uses a load-balancing algorithm to automatically assign end devices in a VSAN to one of the NPV core switch links (in the same VSAN) upon initial login. If there are multiple NPV core switch links in the same VSAN, then you cannot assign a specific one to an end device.
- Both servers and targets can be connected to an NPV device.
- Remote SPAN is not supported.
- Local switching is not supported; all traffic is switched using the NPV core switch.

- NPV devices can connect to multiple NPV core switches. In other words, different NP ports can be connected to different NPV core switches.
- NPV supports NPIV-capable module servers (nested NPIV).
- Only F, NP, and SD ports are supported in NPV mode.
- In the case of servers that are booted over the SAN with NPV, if an NPV link failover occurs, servers will lose access to their boot LUN temporarily.
- NPV switches do not recognize the BB_SCN configuration on the xNP ports because of interoperability issues with the third-party core switches.

NPV Traffic Management Guidelines

When deploying NPV traffic management, follow these guidelines:

- Use NPV traffic management only when the automatic traffic engineering by the NPV device is not sufficient for the network requirements.
- Do not configure traffic maps for all the servers. For non-configured servers, NPV will use automatic traffic engineering.
- Configure the Persistent FC ID on the core switch. Traffic engineering directs the associated server interface to external interfaces that lead to the same core switch. The server will be assigned the same FC ID for every log in. This guideline is not applicable if a 91x4 switch is used as the core switch.
- Server interfaces configured to a set of external interfaces cannot use any other available external interfaces, even if the configured interfaces are not available.
- Do not configure disruptive load balancing because this involves moving a device from one external interface to another interface. Moving the device between external interfaces requires NPV relogin to the core switch through F port leading to traffic disruption.
- Link a set of servers to a core switch by configuring the server to a set of external interfaces that are linked to the core switch.

DPVM Configuration Guidelines

When NPV is enabled, the following requirements must be met before you configure DPVM on the NPV core switch:

- You must explicitly configure the WWN of the internal FLOGI in DPVM. If DPVM is configured on the NPV core switch for an end device that is connected to the NPV device, then that end device must be configured to be in the same VSAN. Logins from a device connected to an NPV device will fail if the device is configured to be in a different VSAN. To avoid VSAN mismatches, ensure that the internal FLOGI VSAN matches the port VSAN of the NP port.
- The first login from an NP port determines the VSAN of that port. If DPVM is configured for this first login, which is the internal login of the NPV device, then the NPV core switch's VSAN F port is located in that VSAN. Otherwise, the port VSAN remains unchanged.

For details about DPVM configuration, see the *Cisco MDS 9000 Family NX-OS Fabric Configuration Guide*.

NPV and Port Security Configuration Guidelines

Port security is enabled on the NPV core switch on a per interface basis. To enable port security on the NPV core switch for devices logging in via NPV, you must adhere to the following requirements:

- The internal FLOGI must be in the port security database so that, the port on the NPV core switch will allow communications and links.
- All of the end device pWWNs must also be in the port security database.

Once these requirements are met, you can enable port security as you would in any other context. For details about enabling port security, see the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Configuring N Port Virtualization

This section includes the following topics:

- Enabling N Port Identifier Virtualization, page 7-10
- Configuring NPV, page 7-10
- Configuring NPV Traffic Management, page 7-12

Enabling N Port Identifier Virtualization

You must globally enable NPIV for all VSANs on the MDS switch to allow the NPIV-enabled applications to use multiple N port identifiers.

Note All of the N port identifiers are allocated in the same VSAN.

Detailed Steps

To enable or disable NPIV on the switch, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# feature npiv`	Enables NPIV for all VSANs on the switch.
Step 3	`switch(config)# no feature npiv`	Disables (default) NPIV on the switch.

Configuring NPV

When you enable NPV, the system configuration is erased and the system reboots with the NPV mode enabled.

Note We recommend that you save the current configuration either on bootflash or a TFTP server before NPV (if the configuration is required for later use). Use the following commands to save either your non-NPV or NPV configuration:

```
switch# copy running bootflash:filename
```

The configuration can be reapplied later using the following command:

```
switch# copy bootflash:filename running-config
```

Detailed Steps

To configure NPV using the CLI, perform the following tasks:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	On the NPV core switch, enters configuration mode.
Step 2	`switch(config)# feature npiv` `switch (config)#`	Enables NPIV mode on the NPV core switch.
	`switch (config)# no feature npiv`	Disables NPIV mode on the NPV core switch.
Step 3	`switch(config)# interface fc 2/1` `switch(config-if)# switchport mode F`	Configures the NPIV core switch port as an F port.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.

Step 4	`switch(config)# vsan database`
	`switch(config-vsan-db)# vsan 8 interface fc 2/1`
	`switch(config-vsan-db)# vsan 10 interface fc 2/1`
		Configures the port VSANs for the F port on the NPIV core switch.

Step 5	`switch(config)# npv enable`	Enables NPV mode on a NPV device (module, Cisco MDS 9124, Cisco MDS 9134, or Cisco MDS 9148 Fabric Switch). The module or switch is rebooted, and when it comes back up, is in NPV mode. **Note** A write-erase is performed during the reboot.
Step 6	`switch(config)# interface fc 1/1` `switch(config-if)# switchport mode NP`	On the NPV device, selects the interfaces that will be connected to the aggregator switch and configure them as NP ports.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.

Step 7	`switch(config)# vsan database`
	`switch(config-vsan-db)# vsan 9 interface fc 1/1`
	`switch(config-vsan-db)# vsan 11 interface fc 1/1`
		Configures the port VSANs for the NP port on the NPV device.

Step 8	`switch(config-if)# exit`	Exits interface mode for the port.
Step 9	`switch(config)# interface fc 1/2 - 6` `switch(config-if)# switchport mode F`	Selects the remaining interfaces (2 through 6) on the NPV-enabled device and configures them as F ports.
	`switch(config-if)# no shutdown`	Changes Admin status to bring up the interfaces.

	Command	Purpose

Step 10	`switch(config)# vsan database`
	`switch(config-vsan-db)# vsan 12 interface fc 1/1 - 6`
	`switch(config-vsan-db)# vsan 13 interface fc 1/1 - 6`
		Configures the port VSANs for the F ports on the NPV device.

Step 11	`switch(config-npv)# no npv enable` `switch(config)#`	Terminates session and disables NPV mode, which results in a reload of the NPV device.

Configuring NPV Traffic Management

The NPV traffic management feature is enabled after configuring NPV. Configuring NPV traffic management involves configuring a list of external interfaces to the servers, and enabling or disabling disruptive load balancing.

This section includes the following topics:

- Configuring List of External Interfaces per Server Interface, page 7-12
- Enabling the Global Policy for Disruptive Load Balancing, page 7-13

Configuring List of External Interfaces per Server Interface

A list of external interfaces are linked to the server interfaces when the server interface is down, or if the specified external interface list includes the external interface already in use.

Detailed Steps

To configure the list of external interfaces per server interface, perform the following tasks:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode on the NPV.
Step 2	`switch(config)# npv traffic-map server-interface` *svr-if-range* `external-interface fc` *ext-fc-if-range* `switch (config)#`	Allows you to configure a list of external FC interfaces per server interface by specifying the external interfaces in the svr-if-range. The server to be linked is specified in the ext-fc-if-range.
	`switch(config)# npv traffic-map server-interface` *svr-if-range* `external-interface port-channel` *ext-pc-if-range* `switch (config)#`	Allows you to configure a list of external PortChannel[1] interfaces per server interface by specifying the external interfaces in the svr-if-range. The server to be linked is specified in the ext-pc-if-range.
	`switch(config)# no npv traffic-map server-interface` *svr-if-range* `external-interface` *ext-if-range* `switch (config)#`	Disables the NPV traffic management feature on the NPV.

1. While mapping non-PortChannel interfaces and PortChannel interfaces to the server interfaces, include them separately in two steps.

Enabling the Global Policy for Disruptive Load Balancing

Disruptive load balancing allows you to review the load on all the external interfaces and balance the load disruptively. Disruptive load balancing is done by moving the servers using heavily loaded external interfaces, to the external interfaces running with fewer loads.

Detailed Steps

To enable or disable the global policy for disruptive load balancing, perform the following tasks:

	Command	Purpose
Step 1	`switch# config t` `switch(config)#`	Enters configuration mode on the NPV.
Step 2	`switch(config)# npv auto-load-balance disruptive` `switch (config)#`	Enables disruptive load balancing on the NPV core switch.
Step 3	`switch (config)# no npv auto-load-balance disruptive`	Disables disruptive load balancing on the NPV core switch.

Verifying NPV Configuration

This section includes the following topics:

- Verifying NPV, page 7-14
- Verifying NPV Traffic Management, page 7-15

To display NPV configuration information, perform one of the following tasks:

Command	Purpose
show fcns database	Displays all the NPV devices in all the VSANs that the aggregator switch belongs to.
show fcns database detail	Displays additional details such as IP addresses, switch names, interface names about the NPV devices.
show npv flogi-table	Displays a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs.
show npv status	Displays the status of the different servers and external interfaces.
show npv traffic-map	Displays the NPV traffic map.
show npv internal info traffic-map	Displays the NPV internal traffic details.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

Verifying NPV

To view all the NPV devices in all the VSANs that the aggregator switch belongs to, enter the **show fcns database** command.

```
switch# show fcns database

VSAN 1:
--------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
--------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init

Total number of entries = 4
```

For additional details (such as IP addresses, switch names, interface names) about the NPV devices you see in the **show fcns database** output, enter the **show fcns database detail** command.

```
switch# show fcns database detail

------------------------
VSAN:1 FCID:0x010000
------------------------
port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/1
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:01:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:01:00:0d:ec:2f:c1:40 (Cisco)

------------------------
VSAN:1 FCID:0x010001
------------------------
port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
node-wwn :20:00:00:0d:ec:2f:c1:40
class :2,3
node-ip-addr :172.20.150.38
ipa :ff ff ff ff ff ff ff ff
fc4-types:fc4_features :npv
symbolic-port-name :para-3:fc1/2
symbolic-node-name :para-3
port-type :N
port-ip-addr :0.0.0.0
fabric-port-wwn :20:02:00:0d:ec:04:99:40
hard-addr :0x000000
permanent-port-wwn (vendor) :20:02:00:0d:ec:2f:c1:40 (Cisco)
```

If you need to contact support, enter the **show tech-support NPV** command and save the output so that support can use it to troubleshoot, if necessary.

To display a list of the NPV devices that are logged in, along with VSANs, source information, pWWNs, and FCIDs, enter the **show npv flogi-table** command.

```
switch# show npv flogi-table
--------------------------------------------------------------------------------
SERVER                                                                  EXTERNAL
INTERFACE VSAN FCID             PORT NAME               NODE NAME       INTERFACE
--------------------------------------------------------------------------------
fc1/19    1    0xee0008 10:00:00:00:c9:60:e4:9a 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee0009 20:00:00:00:0a:00:00:01 20:00:00:00:c9:60:e4:9a fc1/1
fc1/19    1    0xee000a 20:00:00:00:0a:00:00:02 20:00:00:00:c9:60:e4:9a fc1/9
fc1/19    1    0xee000b 33:33:33:33:33:33:33:33 20:00:00:00:c9:60:e4:9a fc1/1

Total number of flogi = 4.
```

To display the status of the different servers and external interfaces, enter the **show npv status** command.

```
switch# show npv status
npiv is enabled

External Interfaces:
====================
  Interface: fc1/1, VSAN: 2, FCID: 0x1c0000, State: Up
  Interface: fc1/2, VSAN: 3, FCID: 0x040000, State: Up

  Number of External Interfaces: 2

Server Interfaces:
==================
  Interface: fc1/7, VSAN: 2, NPIV: No, State: Up
  Interface: fc1/8, VSAN: 3, NPIV: No, State: Up

  Number of Server Interfaces: 2
```

Verifying NPV Traffic Management

To display the NPV traffic map, enter the **show npv traffic-map** command.

```
switch# show npv traffic-map
NPV Traffic Map Information:
----------------------------------------
Server-If       External-If(s)
----------------------------------------
fc1/3           fc1/10,fc1/11
fc1/5           fc1/1,fc1/2
----------------------------------------
```

To display the NPV internal traffic details, enter the **show npv internal info traffic-map** command.

```
switch# show npv internal info traffic-map
NPV Traffic Map Information:
----------------------------------------
Server-If       External-If(s)
----------------------------------------
fc1/3           fc1/10,fc1/11
fc1/5           fc1/1,fc1/2
----------------------------------------
```

Configuring FlexAttach Virtual pWWN

- Information About FlexAttach Virtual pWWN, page 8-1
- Guidelines and Limitations, page 8-3
- Configuring FlexAttach Virtual pWWN, page 8-3
- Verifying FlexAttach Virtual pWWN Configuration, page 8-5
- Monitoring FlexAttach Virtual pWWN, page 8-6

Information About FlexAttach Virtual pWWN

This section includes the following topics:

- FlexAttach Virtual pWWN, page 8-1
- Difference Between San Device Virtualization and FlexAttach Port Virtualization, page 8-2
- FlexAttach Virtual pWWN CFS Distribution, page 8-2
- Security Settings for FlexAttach Virtual pWWN, page 8-3

FlexAttach Virtual pWWN

FlexAttach virtual pWWN feature facilitates server and configuration management. In a SAN environment, the server installation or replacement, requires interaction and coordination among the SAN and server administrators. For coordination, it is important that the SAN configuration does not change when a new server is installed, or when an existing server is replaced. FlexAttach virtual pWWN minimizes the interaction between the server administrator and the SAN administrator by abstracting the real pWWN using virtual pWWNs.

When FlexAttach virtual pWWN is enabled on an interface, a virtual pWWN is assigned to the server interface. The real pWWN is replaced by a virtual pWWN, which is used for a SAN configuration such as zoning.

Server administrators can benefit from FlexAttach in the following scenarios:

- Pre-configure—Pre-configure SAN for new servers that are not available physically yet. For example, they may be on order. FlexAttach can be enabled on the ports designated for the new servers and use the virtual WWNs assigned for configuring SAN. The new servers are then plugged into the fabric without any change needed in the SAN.

- Replacement to the same port—A failed server can be replaced onto the same port without changing the SAN. The new server gets a same pWWN as the failed server because the virtual pWWN is assigned to the port.
- Replacement to (spare)—A spare server, which is on the same NPV device or a different NPV device) can be brought online without changes to the SAN. This action is achieved by moving the virtual port WWN from the current server port to the spare port.
- Server Mobility—A server can be moved to another port on the same NPV device or another NPV device without changing the SAN. This is accomplished by moving the virtual pWWN to the new port. No change is needed if FlexAttach was configured using the physical port WWN of the server to the virtual port WWN mapping.

Difference Between San Device Virtualization and FlexAttach Port Virtualization

Table 8-1 describes the difference between SAN device virtualization (SDV) and FlexAttach port virtualization.

Table 8-1 Difference Between SDV and FlexAttach Virtualization

SAN Device Virtualization (SDV)	FlexAttach Virtualization
Facilitates target and disk management, and only facilitates disk and data migration.	Facilitates server management and has no restriction on the end devices used.
WWN NAT and Fibre Channel ID (FC-ID) are allocated on the virtual device, both primary and secondary.	WWN and Network Address Transport (NAT) is allocated to host bus adapter (HBA).
FC-ID rewrite on the switch indicates a rewrite-capable switch on the path.	No rewrite requirements.
Configuration is distributed. This allows programming rewrites and connectivity anywhere.	Configuration distribution is not required for any of the interface-based configurations.
Configuration is secured to device alias.	Does not require device alias for virtual pWWN.
Does not allow automapping to the secondary device.	Allows automapping to the new HBA. Mapping process is manual for NPIV.

FlexAttach Virtual pWWN CFS Distribution

The FlexAttach virtual pWWN configuration is distributed for CFS through IPv4, and is enabled by default. The FlexAttach virtual pWWN distribution, by default, is on CFS region 201. The CFS region 201 links only to the NPV-enabled switches. Other CFS features such as syslog is on region 0. Region 0 will be linked through IPv4 for all NPV switches on the same physical fabric. If CFS has an option to link through IPv4 or ISL, then CFS will select the ISL path.

Note NPV switches do not have ISL (E or TE ports) and are linked through IPv4.

Security Settings for FlexAttach Virtual pWWN

Security settings for the FlexAttach virtual pWWN feature are done by port security at the NPV core. Node WWN of the end device is used to provide physical security.

For more details on enabling port security, refer to the *Cisco MDS 9000 Family NX-OS Security Configuration Guide*.

Guidelines and Limitations

Following are recommended guidelines and requirements when deploying FlexAttach virtual pWWN:

- FlexAttach configuration is supported only on NPV switches.
- Cisco Fabric Services (CFS) IP version 4 (IPv4) distribution should be enabled.
- Virtual WWNs should be unique across the fabric.

Configuring FlexAttach Virtual pWWN

This section includes the following topics:

- Automatically Assigning FlexAttach Virtual pWWN, page 8-3
- Manually Assigning FlexAttach Virtual pWWN, page 8-4
- Mapping pWWN to Virtual pWWN, page 8-4

Automatically Assigning FlexAttach Virtual pWWN

Automatic assignment of virtual pWWN can be configured on an NPV switch globally, per VSAN, or per port. When assigned automatically, a virtual WWN is generated from the device local switch WWN.

Prerequisites

The port must be in a shut state when the virtual pWWN is enabled.

Detailed Steps

To assign a virtual pWWN automatically, perform this task:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch# (config)# flex-attach virtual-pwwn auto [interface interface-list]`	Assigns FlexAttach virtual pWWN automatically for the interfaces.
	`switch# (config)# flex-attach virtual-pwwn auto [vsan vsan-range]`	Assigns FlexAttach virtual pWWN automatically for the VSANs.
Step 3	`switch# (config)# flex-attach commit`	Commits the configuration.

Troubleshooting Tips

- When the *interface-list* value is not included in the command, virtual pWWN is assigned globally.
- All the interfaces mentioned in the *interface-list* value must be in a shut state.

Manually Assigning FlexAttach Virtual pWWN

You can manually assign a WWN to the interface, without generating it through the switch. Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.

Prerequisites

- Some ports may be in automode, some in manual mode, and the virtual pWWNs need not be assigned.
- The port must be in a shut state when a virtual pWWN is assigned.

Restrictions

- The interface mentioned in the interface value must be in a shut state.

Detailed Steps

To assign virtual pWWN manually, perform this task:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch# (config)# flex-attach virtual-pwwn` *vpwwn* `interface` *interface*	Configures the FlexAttach virtual pWWN for the interface.
	`switch# (config)# flex-attach virtual-pwwn` *vpwwn* `interface` *interface* `[vsan` *vsan]*	Configures the FlexAttach virtual pWWN for the interface in the VSAN.
Step 3	`switch# (config)# flex-attach commit`	Commits the configuration.

Mapping pWWN to Virtual pWWN

You can configure virtual pWWNs through real pWWNs. This process is required for NPIV hosts containing multiple pWWNs, of which only FLOGI is mapped to the virtual pWWN. Subsequent FDSIDs will have different mappings.

Several checks are done by the NPV core to ensure the uniqueness of virtual pWWNs in the switch across the NPV switches. When duplicate virtual pWWNs are configured, the subsequent logins are rejected by the NPV core switch.

Prerequisites

- The interface must be in a shut state and the specified virtual pWWN should not be logged in.

Restrictions

- The specified virtual pWWN and the real pWWN must not be logged in.

Detailed Steps

To map pWWN to virtual pWWN, perform this task:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch# (config)# flex-attach virtual-pwwn` *vpwwn* **pwwn** *pwwn*	Maps the pWWN to the virtual pWWN.
Step 3	`switch# (config)# flex-attach commit`	Commits the configuration.

Verifying FlexAttach Virtual pWWN Configuration

To display FlexAttach configuration information, perform one of the following tasks:

Command	Purpose
show flex-attach virtual-wwn	Displays the type and value of virtual pWWNs.
show fcns database	Displays if the end device is logged with the correct virtual WWNs.

For detailed information about the fields in the output from these commands, refer to the *Cisco MDS NX-OS Command Reference*.

To view and confirm that the type and value of virtual pWWNs are correct, enter the **show flex-attach virtual-wwn** command. (See Example 8-1.)

Example 8-1 Displaying the Type and Value of Virtual pWWNs

```
switch# show flex-attach virtual-wwn
VIRTUAL PORT WWNS ASSIGNED TO INTERFACES
------------------------------------------------------------------------
VSAN     INTERFACE  VIRTUAL-PWWN               AUTO    LAST-CHANGE
------------------------------------------------------------------------
1        fc1/1      00:00:00:00:00:00:00:00
1        fc1/2      22:73:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/3      22:5e:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/4      22:5f:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/5      22:74:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:26:24 2008
1        fc1/6      22:60:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/7      22:61:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/8      22:62:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/9      22:63:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/10     22:64:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1        fc1/11     22:65:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
1         fc1/12     22:66:00:05:30:01:6e:1e   TRUE   Thu Jan 31 01:58:52 2008
```

Verifying the End Device

To verify that the end device is logged with the correct virtual WWNs, use the **show fcns database** command on the NPV core. (See Example 8-2.)

Example 8-2 Verifying the End Device

```
switch# show fcns database
VSAN 1:
------------------------------------------------------------------------
FCID TYPE PWWN (VENDOR) FC4-TYPE:FEATURE
------------------------------------------------------------------------
0x010000 N 20:01:00:0d:ec:2f:c1:40 (Cisco) npv
0x010001 N 20:02:00:0d:ec:2f:c1:40 (Cisco) npv
0x010200 N 21:00:00:e0:8b:83:01:a1 (Qlogic) scsi-fcp:init
0x010300 N 21:01:00:e0:8b:32:1a:8b (Qlogic) scsi-fcp:init
Total number of entries = 4
```

Monitoring FlexAttach Virtual pWWN

Table 8-2 lists the errors that might be displayed and provides the workarounds.

Table 8-2 FlexAttach Errors and Workarounds

Error	Description	Workaround
fc1/1 : interface is not down	FlexAttach configuration fails because the configuration is enabled for an active interface with the operation state as up.	To move the port to the shut state, enable the FlexAttach configuration, and then move the port to no shut state.
FlexAttach configuration is not distributed to the peers	The FlexAttach configuration on one peer NPV is not available to any other peer NPV.	FlexAttach configuration will not be distributed if **cfs ipv4 distribute**, or **cfs ipv6 distribute** is disabled. Enable **cfs ipv4 distribute**, or **cfs ipv6 distribute**.
Even with CFS distribution enabled Inagua does not become a peer with other NPVs	CFS over IP is enabled, and the Inagua in one blade center is not the peer NPV for other NPVs.	CFS over IP uses IP multicast to discover the NPV peers in the network. IBM MM does not support multicast and cannot act as a peer with NPV. This prevents the FlexAttach configuration from getting distributed to other peer NPVs in the network.
NP port uses physical pWWN instead of virtual pWWN configured through FlexAttach	This occurs when NP port uses physical pWWN instead of virtual pWWN, that is configured through FlexAttach.	FlexAttach is supported on server interfaces such as F ports, and not on external interfaces such as NP ports.

Error	Description	Workaround
real port WWN and virtual WWN cannot be same	This occurs when you try to configure FlexAttach with a similar value for pWWN and virtual pWWN.	Use different values for pWWN and virtual pWWN, as similar values for pWWN and virtual pWWn are not allowed.
Virtual port WWN already exists	This occurs when you try to configure an already defined pWWN to a different interface.	Use an undefined virtual pWWN for a new interface.

Configuring Port Tracking

The port tracking feature is unique to the Cisco MDS 9000 Family of switches. This feature uses information about the operational state of the link to initiate a failure in the link that connects the edge device. This process of converting the indirect failure to a direct failure triggers a faster recovery process towards redundant links. When enabled, the port tracking feature brings down the configured links based on the failed link and forces the traffic to be redirected to another redundant link.

This chapter includes the following sections:

- Information About Port Tracking, page 9-1
- Guidelines and Limitations, page 9-2
- Default Settings, page 9-2
- Configuring Port Tracking, page 9-3Configuring Port Tracking, page 9-3
- Displaying Port Tracking Information, page 9-6

Information About Port Tracking

Generally, hosts can instantly recover from a link failure on a link that is immediately (direct link) connected to a switch. However, recovering from an indirect link failure between switches in a WAN or MAN fabric with a keep-alive mechanism is dependent on several factors such as the time out values (TOVs) and on registered state change notification (RSCN) information.

In Figure 9-1, when the direct link 1 to the host fails, recovery can be immediate. However, when the ISL 2 fails between the two switches, recovery depends on TOVs, RSCNs, and other factors.

Figure 9-1 Traffic Recovery Using Port Tracking

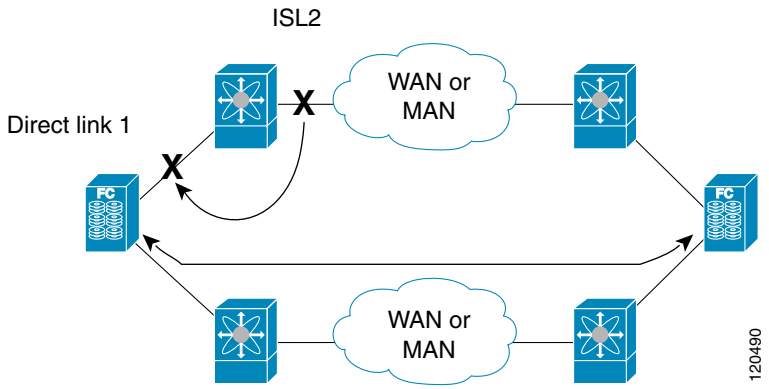

The port tracking feature monitors and detects failures that cause topology changes and brings down the links connecting the attached devices. When you enable this feature and explicitly configure the linked and tracked ports, the Cisco NX-OS software monitors the tracked ports and alters the operational state of the linked ports on detecting a link state change.

The following terms are used in this chapter:

- Tracked ports—A port whose operational state is continuously monitored. The operational state of the tracked port is used to alter the operational state of one or more ports. Fibre Channel, VSAN, PortChannel, FCIP, or a Gigabit Ethernet port can be tracked. Generally, ports in E and TE port modes can also be Fx ports.
- Linked ports—A port whose operational state is altered based on the operational state of the tracked ports. Only a Fibre Channel port can be linked.

Guidelines and Limitations

Before configuring port tracking, consider the following guidelines:

- Verify that the tracked ports and the linked ports are on the same Cisco MDS switch.
- Do not track a linked port back to itself (for example, Port fc1/2 to Port fc2/5 and back to Port fc1/2) to avoid recursive dependency.
- Be aware that the linked port is automatically brought down when the tracked port goes down. Be aware that the linked port is automatically brought down when the tracked port goes down.

Default Settings

Table 9-1 lists the default settings for port tracking parameters.

Table 9-1 Default Port Tracking Parameters

Parameters	Default
Port tracking	Disabled.
Operational binding	Enabled along with port tracking.

Configuring Port Tracking

Port tracking has the following features:

- The application brings the linked port down when the tracked port goes down. When the tracked port recovers from the failure and comes back up again, the tracked port is also brought up automatically (unless otherwise configured).
- You can forcefully continue to keep the linked port down, even though the tracked port comes back up. In this case, you must explicitly bring the port up when required.

This section includes the following topics:

- Enabling Port Tracking, page 9-3
- Information About Configuring Linked Ports, page 9-3
- Binding a Tracked Port Operationally, page 9-4
- Information About Tracking Multiple Ports, page 9-4
- Tracking Multiple Ports, page 9-5
- Information About Monitoring Ports in a VSAN, page 9-5
- Monitoring Ports in a VSAN, page 9-5
- Information AboutForceful Shutdown, page 9-6
- Forcefully Shutting Down a Tracked Port, page 9-6

Enabling Port Tracking

The port tracking feature is disabled by default in all switches in the Cisco 9000 Family. When you enable this feature, port tracking is globally enabled for the entire switch.

To configure port tracking, enable the port tracking feature and configure the linked port(s) for the tracked port.

Detailed Steps

To enable port tracking, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# feature port-track`	Enables port tracking.
	`switch(config)# no feature port-track`	Removes the currently applied port tracking configuration and disables port tracking.

Information About Configuring Linked Ports

You can link ports using one of two methods:

- Operationally binding the linked port(s) to the tracked port (default).
- Continuing to keep the linked port down forcefully—even if the tracked port has recovered from the link failure.

Binding a Tracked Port Operationally

When you configure the first tracked port, operational binding is automatically in effect. When you use this method, you have the option to monitor multiple ports or monitor ports in one VSAN.

Detailed Steps

To operationally bind a tracked port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6` `switch(config-if)#`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports. **Note** This link symbolizes the direct link (1) in Figure 9-1.
Step 3	`switch(config-if)# port-track interface port-channel 1`	Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (2) in Figure 9-1.
	`switch(config-if)# no port-track interface port-channel 1`	Removes the port tracking configuration that is currently applied to interface fc8/6.

Information About Tracking Multiple Ports

You can control the operational state of the linked port based on the operational states of multiple tracked ports. When more than one tracked port is associated with a linked port, the operational state of the linked port will be set to down only if all the associated tracked ports are down. Even if one tracked port is up, the linked port will stay up.

In Figure 9-2, only if both ISLs 2 and 3 fail, will the direct link 1 be brought down. Direct link 1 will not be brought down if either 2 or 3 are still functioning as desired.

Figure 9-2 Traffic Recovery Using Port Tracking

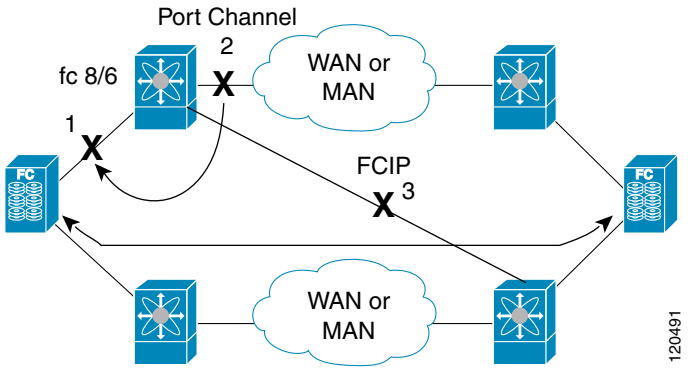

Tracking Multiple Ports

Detailed Steps

To track multiple ports, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports. **Note** This link symbolizes the direct link (1) in Figure 9-2.
Step 3	`switch(config-if)# port-track interface port-channel 1`	Tracks interface fc8/6 with interface port-channel 1. When port-channel 1 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (2) in Figure 9-2.
Step 4	`switch(config-if)# port-track interface fcip 5`	Tracks interface fc8/6 with interface fcip 5. When FCIP 5 goes down, interface fc8/6 is also brought down. **Note** This link symbolizes the ISL (3) in Figure 9-2.

Information About Monitoring Ports in a VSAN

You can optionally configure one VSAN from the set of all operational VSANs on the tracked port with the linked port by specifying the required VSAN. This level of flexibility provides higher granularity in tracked ports. In some cases, when a tracked port is a TE port, the set of operational VSANs on the port can change dynamically without bringing down the operational state of the port. In such cases, the port VSAN of the linked port can be monitored on the set of operational VSANs on the tracked port.

If you configure this feature, the linked port is up only when the VSAN is up on the tracked port.

Tip The specified VSAN does not have to be the same as the port VSAN of the linked port.

Monitoring Ports in a VSAN

Detailed Steps

To monitor a tracked port in a specific VSAN, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc8/6`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.
Step 3	`switch(config-if)# port-track interface port-channel 1 vsan 2`	Enables tracking of the PortChannel in VSAN 2.
	`switch(config-if)# no port-track interface port-channel 1 vsan 2`	Removes the VSAN association for the linked port. The PortChannel link remains in effect.

Information AboutForceful Shutdown

If a tracked port flaps frequently, then tracking ports using the operational binding feature may cause frequent topology change. In this case, you may choose to keep the port in the down state until you are able to resolve the reason for these frequent flaps. Keeping the flapping port in the down state forces the traffic to flow through the redundant path until the primary tracked port problems are resolved. When the problems are resolved and the tracked port is back up, you can explicitly enable the interface.

Tip If you configure this feature, the linked port continues to remain in the shutdown state even after the tracked port comes back up. You must explicitly remove the forced shut state (by administratively bringing up this interface) of the linked port once the tracked port is up and stable.

Forcefully Shutting Down a Tracked Port

Detailed Steps

To forcefully shut down a tracked port, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/5`	Configures the specified interface and enters the interface configuration submode. You can now configure tracked ports.
Step 3	`switch(config-if)# port-track force-shut`	Forcefully shuts down the tracked port.
	`switch(config-if)# no port-track force-shut`	Removes the port shutdown configuration for the tracked port.

Displaying Port Tracking Information

Examples

The **show** commands display the current port tracking settings for the Cisco MDS switch (see Examples 9-1 to 9-4).

Example 9-1 Displays the Linked and Tracked Port Configuration

```
switch# show interface
...
fc8/6 is down (All tracked ports down)     <-----------------------------------Linked port
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 21:c6:00:05:30:00:37:1e
    Admin port mode is auto, trunk mode is on
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface port-channel 1 vsan 2 (trunking) <-----Tracked port
    Port tracked with interface fcip 5 <-----------------------------------Tracked port
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
```

```
      269946 frames input, 22335204 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      205007 frames output, 10250904 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      2 output OLS, 2 LRR, 0 NOS, 1 loop inits
      0 receive B2B credit remaining
      0 transmit B2B credit remaining
...
```

Example 9-2 Displays a Tracked Port Configuration for a Fibre Channel Interface

```
switch# show interface fc1/1
fc1/1 is down (Administratively down)
    Hardware is Fibre Channel, FCOT is short wave laser w/o OFC (SN)
    Port WWN is 20:01:00:05:30:00:0d:de
    Admin port mode is FX
    Port vsan is 1
    Receive data field Size is 2112
    Beacon is turned off
    Port tracked with interface fc1/2 (down)
    Port tracked with interface port-channel 1 vsan 2 (down)
    Port tracked with interface fcip1 (down)
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      1 frames input, 128 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      1 frames output, 128 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
      0 receive B2B credit remaining
      0 transmit B2B credit remaining
```

Example 9-3 Displays a Tracked Port Configuration for a PortChannel Interface

```
switch# show interface port-channel 1
port-channel 1 is down (No operational members)
    Hardware is Fibre Channel
    Port WWN is 24:01:00:05:30:00:0d:de
    Admin port mode is auto, trunk mode is on
    Port vsan is 2
    Linked to 1 port(s)
      Port linked to interface fc1/1
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      0 frames input, 0 bytes
        0 discards, 0 errors
        0 CRC,  0 unknown class
        0 too long, 0 too short
      0 frames output, 0 bytes
        0 discards, 0 errors
      0 input OLS, 0 LRR, 0 NOS, 0 loop inits
      0 output OLS, 0 LRR, 0 NOS, 0 loop inits
    No members
```

Example 9-4 Displays a Forced Shutdown Configuration

```
switch# show interface fc 1/5

fc1/5 is up
    Hardware is Fibre Channel, FCOT is short wave laser
    Port WWN is 20:05:00:05:30:00:47:9e
    Admin port mode is F
    Port mode is F, FCID is 0x710005
    Port vsan is 1
    Speed is 1 Gbps
    Transmit B2B Credit is 64
    Receive B2B Credit is 16
    Receive data field Size is 2112
    Beacon is turned off
    Port track mode is force_shut <--this port remains shut even if the tracked port is back up
```

Symbols

* (asterisk)
  first operational port 6-20

Numerics

12-port 4-Gbps switching modules
  BB_credit buffers 4-13
  configuration guidelines 3-28
  default settings 3-25
  See also switching modules
16-port switching modules
  configuring BB_credits 4-1
  LEDs 2-11
  See also switching modules
24-port 4-Gbps switching modules
  bandwidth fairness 3-17
  default settings 3-25
  example configurations 3-46, 3-49, 4-12
  oversubscription 3-11
  shared resources 3-11
  See also switching modules
24-port 8-Gbps switching modules
  default settings 3-25
  example configurations 4-8
32-port switching modules
  configuring BB_credits 4-1
  See also switching modules
4/44-port 8-Gbps switching modules
  default settings 3-25
  example configurations 4-9
48-port 4-Gbps switching modules
  bandwidth fairness 3-17
  default settings 3-25
  example configurations 3-47, 3-48, 4-10
  oversubscription 3-11
  shared resources 3-11
  See also switching modules
48-port 8-Gbps switching modules
  default settings 3-25
  example configurations 3-46, 4-7
  See also switching modules
4-port 10-Gbps switching modules
  BB_credit buffers 4-15
  configuration guidelines 3-28
  default settings 3-25
  See also switching modules

A

ACL adjacency sharing
  disabling for downgrading 3-43
administrative speeds
  configuring 2-10, 2-25
administrative states
  description 2-6
  setting 2-22
ALPA caches
  allocation 2-50
  clearing 2-29
  description 2-14
  displaying contents 2-51
  inserting entries manually 2-29
auto mode
  configuring 2-22

auto port mode
description 2-5
interface configuration 2-2
autosensing speed
Generation 2 switching modules 2-10

B

bandwidth fairness
disabling 3-41
enabling 3-40
BB_credit buffers
12-port 4-Gbps switching module allocations 4-13
12-port 4-Gbps switching module considerations 4-14
24-port 4-Gbps switching module allocations 4-12
24-port 4-Gbps switching module considerations 4-12, 4-13
24-port 8-Gbps switching module considerations 4-8
4/44-port 8-Gbps switching module considerations 4-9
48-port 4-Gbps switching module considerations 4-10
48-port 8-Gbps switching module considerations 4-7
4-port 10-Gbps switching module allocations 4-15
4-port 10-Gbps switching module considerations 4-15, 4-16
allocation defaults (table) 4-6, 4-7, 4-8, 4-9, 4-10
BB_credits
configuring 4-21
description 4-1
reason codes 2-8
BB_SC
enabling 4-20, 4-24
beacon modes
configuring 2-27
identifying LEDs 2-11
bit errors
reasons 2-12
bit error thresholds
configuring 2-11, 2-28
description 2-11, 2-28
B port mode
description 2-5
interface modes 2-5
bridge port mode. See B port mode
buffer pools
Generation 2 switching modules 4-2
buffer-to-buffer credits. See BB_credits
buffer-to-buffer start change. See BB_SC

C

Cisco MDS 9216i switches
configuring extended BB_credits 4-18
configuring NPV 7-10

D

dedicated rate mode
description 3-7
migrating from shared rate mode 3-26
migrating to shared rate mode 3-27
destination IDs
exchange based 6-5
flow based 6-4
domain IDs
assignment failures 2-9
domain manager
isolation 2-9
downgrading
disabling ACL adjacency sharing 3-43
dynamic bandwidth management
description 3-10

E

EISLs

PortChannel links 6-2
enhanced ISLs. See EISLs
E port mode
classes of service 2-3
description 2-3
E ports
32-port guidelines 2-18
32-port switching module configuration guidelines 6-14
configuring 2-22
isolation 2-9
exchange link parameter. See ELP
expansion port mode. See E port mode
extended BB_credits
configuring 4-23
displaying information 4-25
Generation 2 switching modules 4-17
licensing 4-17

F

fabric loop port mode. See FL port mode
fabric port mode. See F port mode
fcdomains
overlap isolation 2-9
Fiber Channel interfaces
configuring system default port mode f 2-23
Fibre Channel interfaces
administrative states 2-6
BB_credits 4-1
configuring 2-21
configuring auto port mode 2-22
configuring beacon modes 2-27
configuring bit error thresholds 2-11, 2-28
configuring descriptions 2-26
configuring frame encapsulation 2-10
configuring port modes 2-22
configuring port speed group 2-26
configuring receive data field sizes 4-21
configuring speeds 2-10, 2-25
default settings 2-20, 3-25, 5-11, 6-16
deleting from PortChannels 6-9
disabling 2-22
displaying capabilities on Generation 2 switching modules 3-44
displaying information 2-42 to 2-49
enabling 2-22
graceful shutdown 2-9
modes 2-2 to 2-6
operational states 2-6
performance buffers 4-2
reason codes 2-6
states 2-6
taking out of service on Generation 2 switching modules 3-41
troubleshooting operational states 2-8
See also interfaces 2-6
FL port mode
classes of service 2-3
description 2-3
FL ports
configuring 2-22
description 2-3
nonparticipating code 2-9
See also Fx ports
F port mode
classes of service 2-3
description 2-3
F ports
configuring 2-22
description 2-3
See also Fx ports
frame encapsulation
configuring 2-10
Fx ports
32-port default 2-19
configuring 2-22
description 2-5

interface modes 2-5
See also F ports; FL ports 2-5

G

Generation 1 switching modules
extended BB_credits 4-18
port index allocations 3-19
Generation 2 switching modules
buffer groups 4-2 to 4-15
configuring port speeds 3-29
configuring rate modes 3-30
default settings 3-25
description 3-1 to 3-3
disabling ACL adjacency sharing 3-43
displaying port resources 3-42
dynamic bandwidth management 3-10
extended BB_credits 4-17, 4-19
interface capabilities 3-44
out-of-service interfaces 3-11
port groups 3-3
port index allocations 3-19
port rate modes 3-5
recovering from powered-down state 3-21
releasing shared resources 3-42
taking interfaces out of service 3-41
Generation 3 switching modules
default settings 3-25

I

indirect link failures
recovering 9-1
interfaces
adding to PortChannels 6-8, 6-18
configuring descriptions 2-26
default settings 2-20, 3-25, 5-11, 6-16
deleting from PortChannels 6-9
displaying information 2-42 to 2-49
displaying SFP information 2-48
forced addition to PortChannels 6-9, 6-19
isolated states 6-9
suspended states 6-9
IPv4 default gateways
configuring mgmt0 interfaces 2-17
ISLs
PortChannel links 6-2

L

LEDs
beacon mode states 2-11
speed 2-11
licenses
extended BB_credits 4-17, 4-18
link failures
recovering 9-1
load balancing
description 6-4
PortChannels 6-2

M

management interfaces
configuring 2-17, 2-37
default settings 2-20, 3-25, 5-11, 6-16
displaying information 2-55
features 2-17
mgmt0 interfaces
configuring 2-17, 2-37
default settings 2-20, 3-25, 5-11, 6-16
features 2-17
MPS-14/2 modules
configuring extended BB_credits 4-18

N

NL ports
- interface modes 2-6

nonparticipating codes
- description 2-9

NPIV
- description 7-2
- enabling 7-10

NP links 7-5

N port identifier virtualization. See NPIV

NL ports
- See also Nx ports

NP-ports 7-5

NPV, configuring 7-10

NPV mode 7-4

O

operational states
- configuring on Fibre Channel interfaces 2-22
- description 2-6

out-of-service interfaces
- description 3-11

oversubscription
- disabling restrictions 3-37
- enabling restrictions 3-39
- Generation 2 switching modules 3-11
- ratios 3-12

P

performance buffers
- configuring 4-22
- description 4-2

PortChannel modes
- description 6-6

PortChannel Protocol
- autocreation 6-11
- configuring autocreation 6-20
- converting autocreated groups to manually configured 6-12
- creating channel group 6-10
- description 6-9
- enabling autocreation 6-20

PortChannels
- adding interfaces 6-8, 6-18
- administratively down 2-9
- comparison with trunking 6-3
- compatibility checks 6-8
- creating 6-17
- default settings 6-16
- deleting 6-7
- deleting interfaces 6-9
- examples 6-2
- forcing interface additions 6-9, 6-19
- Generation 2 switching module interfaces 3-21
- interface states 6-9
- load balancing 6-4
- misconfiguration error detection 6-13
- verifying configurations 6-21 to 6-24

port groups
- assigning extended BB_credits 4-18
- description 3-3
- Generation 2 Fibre Channel switching modules 3-3
- Generation 3 Fibre Channel switching modules 3-9

port indexes
- description 3-19

port modes
- auto 2-5
- description 2-2 to 2-6

port rate modes
- configuring 3-30
- dedicated 3-7
- description 3-5
- oversubscribed 3-8
- shared 3-8
- See also rate modes

port speed group
configuring 2-26
port speeds
configuring 2-10, 2-25
configuring on Generation 2 switching module interfaces 3-29
displaying configuration 3-30
port tracking
default settings 9-2
description 9-1
displaying information 9-6
enabling 9-3
guidelines 9-2
monitoring ports in a VSAN 9-5
multiple ports 9-4
shutting down ports forcefully 9-6
private devices
TL ports 2-50

R

rate modes
configuring on Generation 2 switching module interfaces 3-30
verifying configuration 3-31
See also port rate modes
reason codes
description 2-6
receive buffer groups. See buffer groups
receive data field sizes
configuring 4-21
recovery
from powered-down state 3-21

S

SD port mode
description 2-5
interface modes 2-5
SD ports
configuring 2-22
SFPs
displaying transmitter types 2-48
transmitter types 2-12
shared rate mode
description 3-8
migrating from dedicated rate mode 3-27
migrating to dedicated rate mode 3-26
oversubscription 3-11
source IDs
exchange based 6-5
flow based 6-4
SPAN destination port mode. See SD port mode
SPAN tunnel port mode. See ST port mode
ST port mode
description 2-5
interface modes 2-5
limitations 2-5
ST ports
interface modes 2-5
subnet masks
configuring mgmt0 interfaces 2-17
switch ports
configuring attribute default values 2-28

T

TE port mode
classes of service 2-4
description 2-4
TE ports
trunking restrictions 5-3
TF port mode
classes of service 2-4
description 2-4
TL port mode
classes of service 2-4
description 2-4

TL ports
ALPA caches 2-14
configuring 2-22, 2-29
description 2-13
displaying information 2-50
private devices 2-50
virtual devices 2-51
tracked ports
binding operationally 9-4
translative loop port mode. See TL port mode
trunk-allowed VSAN lists
description 5-5 to 5-7
trunking
comparison with PortChannels 6-3
configuration guidelines 5-10
configuring modes 5-4
link state 5-4
merging traffic 5-10
restrictions 5-3, 5-7
trunking E port mode. See TE port mode
trunking F port mode. See TF port mode
trunking protocol
default settings 5-11
default state 5-4
detecting port isolation 5-10
disabling 5-11, 5-12
enabling 5-11, 5-12
trunk mode
administrative default 2-29
configuring 5-4, 5-12
status 5-4
trunk ports
displaying information 5-14

V

virtual devices
TL ports 2-51
VSAN IDs
allowed list 5-11
multiplexing traffic 2-4
VSAN interfaces
creating 2-38
description 2-18
displaying information 2-55
VSANs
allowed-active 5-7
configuring allowed-active lists 5-12
configuring trunk-allowed lists 5-5 to ??
loop devices 2-50
mismatches 2-9
port tracking 9-5
TE port mode 2-4
TF port mode 2-4
trunk-allowed 5-7
VSAN trunking. See trunking

W

WWNs
suspended connections 2-9

Z

zones
merge failures 2-9